Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 29 March 2016 (MN LATEX style file v2.2)

Cosmic shear measurement with maximum likelihood and maximum a posteriori inference

Alex Hall⋆ and Andy Taylor
Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh, EH9 3HJ, U.K.

arXiv:1603.08431v1 [astro-ph.CO] 28 Mar 2016

29 March 2016

ABSTRACT

We investigate the problem of noise bias in maximum likelihood and maximum a posteriori estimators for cosmic shear. We derive the leading and next-to-leading order biases and compute them in the context of galaxy ellipticity measurements, extending previous work on maximum likelihood inference for weak lensing. We show that a large part of the bias on these point estimators can be removed using information already contained in the likelihood when a galaxy model is specified, without the need for external calibration. We test these bias-corrected estimators on simulated galaxy images similar to those expected from planned space-based weak lensing surveys, with very promising results. We find that the introduction of an intrinsic shape prior mitigates noise bias, such that the maximum a posteriori estimate can be made less biased than the maximum likelihood estimate. Second-order terms offer a check on the convergence of the estimators, but are largely sub-dominant. We show how biases propagate to shear estimates, demonstrating in our simple setup that shear biases can be reduced by orders of magnitude and potentially to within the requirements of planned space-based surveys. We find that second-order terms can exhibit significant cancellations at low signal-to-noise when Gaussian noise is assumed, which has implications for inferring the performance of shear-measurement algorithms from simplified simulations. We discuss the viability of our point estimators as tools for lensing inference, arguing that they allow for the robust measurement of ellipticity and shear.

Key words: cosmology: theory - gravitational lensing: weak - cosmology: observations - methods: statistical

1 INTRODUCTION

Weak gravitational lensing is developing into a powerful observational probe of the dark matter on large scales. Current cosmic shear surveys are able to break parameter degeneracies affecting more established observables (Heymans et al. 2012; Planck Collaboration et al. 2015) and current galaxy-galaxy lensing surveys provide competitive constraints on the cosmological parameters (Mandelbaum et al. 2013). However, if weak lensing is to fulfil its potential of providing decisive constraints on the geometry of the large-scale Universe, a number of contaminating systematic errors will have to be resolved.

One particularly troublesome systematic arises when the fundamental lensing observable, the shear field, is estimated. In this work we will be concerned with a bias that arises in the shear estimator due to the pixel noise in observed galaxy images, sometimes referred to as 'noise bias' (Bernstein & Jarvis 2002; Hirata & Seljak 2003). This arises due to the non-linear mapping from the observed raw photon count in CCD pixels to the shear map. The former contains shot noise from the discreteness of the counts, as well as a fluctuating sky background and read-out noise from the CCD itself (Mandelbaum et al. 2014). Point estimators for galaxy shape parameters will necessarily be biased by the non-linear transformation of this noise. This is potentially disastrous if, as is usual practice, shear is estimated by averaging ellipticity estimates over sources (Refregier et al. 2012), since the error bar on the shear estimate will shrink with the number of sources whilst the bias will not.

Shear estimation algorithms therefore generally need to be calibrated against image simulations to correct for these biases, which is clearly not an ideal situation since it relies on the existence of simulations specific to a given survey against which to calibrate (Bartelmann et al. 2012). These simulations must be accurate enough and large enough to characterise performance of an algorithm to allow interpolation to real data. In addition, if the simulations themselves change the algorithm must be re-calibrated. Of course, if the bias was known with infinite accuracy it could be sub-

⋆ ahall@roe.ac.uk

© 0000 RAS

tracted off with no loss of precision. The uncertainty enters when simulations are needed for calibration.

It therefore seems desirable that a method is found requiring minimal external calibration for bias. In this work, we investigate how the method of Maximum Likelihood (ML) might be adapted for use as an unbiased shear estimator, by subtracting off from the Maximum-Likelihood Estimate (MLE) a bias predicted from the likelihood, without external calibration. The MLE is a natural choice for estimating galaxy shape parameters, as it is easy to compute and possesses nice asymptotic properties such as consistency, normality, and efficiency. The MLE has already been implemented as a shear estimator with external bias calibration, with good results (Gentile et al. 2012; Zuntz et al. 2013; Mandelbaum et al. 2015; Jarvis et al. 2015). Since the noise bias scales roughly as the inverse-square of the signal-to-noise ratio (S/N), we are forced to consider the next-to-leading-order behaviour of the MLE around its asymptotic distribution, a Gaussian centred on the truth with $\mathcal{O}(\sqrt{n})$ fluctuations, where n is the number of independent data points. It turns out that the MLE is still an efficient estimator at this order in n, a property that may be shown to hold at yet higher order still (Rao 1961). Thus, the MLE corrected for noise bias seems like a promising choice as a point-estimator for shape and shear, both in terms of its bias and mean-squared error. Similar statements may be made regarding the Maximum A Posteriori estimate (MAP), which arises when a prior is multiplied against the likelihood and the mode of the resulting distribution is found. In this work we also investigate this hitherto unexplored shear estimator.

A further advantage of the MLE (or MAP) is that it is naturally complementary to a Bayesian analysis pipeline for weak lensing. The Bayesian approach to inference can avoid the biases inherent to point estimators (Schneider et al. 2015), and is a preferable approach to weak lensing if it can be made computationally feasible. As the likelihood function is a necessary component of a Bayesian hierarchical model, the MLE can be easily incorporated into such a scheme.

This work will deal with the fundamentals of the MLE and the MAP as lensing estimators, using simple models to elucidate the key properties. We will be particularly concerned with how shear estimates respond to ellipticity bias, and how ellipticity bias is affected by the pixellization of a galaxy image. Our results are not specific to any particular lensing survey, but as a reference we present out final bias measurements alongside the total requirements of a fiducial space-based lensing survey. The stable point-spread function (PSF) of space-based lensing surveys such as *Euclid* and WFIRST potentially means that other systematics, including noise bias, might become limiting factors in the inference of lensing power spectra and cosmological parameters. We also expect our results to be of use for current and upcoming ground-based lensing surveys such as the Kilo-Degree Survey (Kuijken et al. 2015), the Dark Energy Survey (Becker et al. 2015), or the Large Synoptic Survey Telescope (Chang et al. 2013). Looking beyond weak lensing, many of our results are of more general applicability, concerning as they do the properties of bias in the MLE and the MAP.

In Section 2 we review the calculation of the MLE noise bias, before specialising to the case of estimating ellipticity from galaxy images. In Section 3 we test our bias-correction formalism on simulated noisy galaxy images, and in Section 4 introduce the MAP as an alternative point estimator. In Section 5 we consider how the bias propagates to one-dimensional shear estimates, extending this to two dimensions in Section 6. In Section 7 we compare our formalism to other methods in the literature, and conclude.

2 NOISE BIAS IN THE MLE

In this section we review the derivation of bias in the MLE, following McCullagh (1987). For simplicity we specialise to a one-dimensional parameter space, but the generalisation to arbitrary dimension is straightforward. We do not assume anything about the form of the likelihood function except that it is differentiable in a sufficiently small region around the true parameter value, making this derivation more general than that presented in the context of weak lensing in Refregier et al. (2012).

Let the log-likelihood be $L(\beta;\mathbf{Y}) \equiv \log p(\mathbf{Y}|\beta)$, where $\mathbf{Y}$ is a data vector and β is the true (but unknown) value of the parameter. Suppose we make n observations, which result in n realisations of $\mathbf{Y}$, which we label $\mathbf{Y}_i$ with $i = 1...n$. We will assume that the $\mathbf{Y}_i$ are independent but not necessarily identically distributed, each drawn from the distribution $p_i(\mathbf{Y}_i|\beta)$. The joint log-likelihood is then

$$L(\beta) \equiv L(\beta;\mathbf{Y}) = \sum_{i=1}^{n} \log p_i(\mathbf{Y}_i|\beta). \quad (1)$$

The MLE (denoted $\hat{\beta}$) is defined implicitly by the equation $L'(\hat{\beta}) = 0$, with a prime denoting differentiation with respect to β.

Now, recall the Central Limit Theorem (CLT). Suppose that X_j are a set of independent (but not necessarily identically distributed) vector-valued random variables, with j labelling members of the set, each with finite mean and variance. The central limit theorem then holds that for sufficiently large n

$$\frac{1}{\sqrt{n}} \sum_{j=1}^{n} (X_j^i - \mu_j^i) \sim N\left(0, \frac{1}{n}\sum_{j=1}^{n} \mathbf{C}_j\right), \quad (2)$$

where μ_j^i is the i-th component of the mean of the j-th vector, and $\mathbf{C}_j$ is the covariance matrix of the j-th vector. Note that we have glossed over some formal regularity conditions.

Now, expand the ML equation around β:

$$L'(\hat{\beta}) \approx L'(\beta) + (\hat{\beta}-\beta)L''(\beta) + \frac{1}{2}(\hat{\beta}-\beta)^2 L'''(\beta) = 0, \quad (3)$$

where we have neglected higher-order terms. Note that this makes sense in the frequentist approach to statistics since β is not considered to be a random variable. Our aim is to find the asymptotic form of this equation in the limit of large n. We know that MLEs are asymptotically normal from the CLT, such that in the limit of large n

$$\hat{\delta} \equiv \sqrt{n}(\hat{\beta}-\beta) \sim N(0, F^{-1}), \quad (4)$$

where the total Fisher information, F, is given by

$$F = -\sum_{j=1}^{n} \left\langle \frac{\partial^2}{\partial\beta^2} \log p_j(\mathbf{Y}_j|\beta) \right\rangle, \quad (5)$$

with angle brackets denoting an average over realisations of $\mathbf{Y}_i$. Now, from the definition of the log-likelihood

$$L'(\beta) = \sum_{j=1}^{n} \frac{\partial}{\partial\beta} \log p_j(\mathbf{Y}_j|\beta). \quad (6)$$

The mean of each term in the above sum is zero, so by the CLT

$$n^{-1/2} L'(\beta) \sim N(0, \bar{\mathbf{C}}), \quad (7)$$

where $\bar{\mathbf{C}}$ is the covariance matrix of the data averaged over n realisations. In other words, $L'(\beta)$ is $\mathcal{O}(\sqrt{n})$, since the average covariance matrix is $\mathcal{O}(1)$. The $\mathcal{O}$ notation here refers to the magnitude of typical r.m.s. fluctuations in a quantity. In this notation, cumulants of individual measurements are taken to be $\mathcal{O}(1)$. Thus we can define a scaled quantity $\mathcal{Z}^{(1)}(\beta) = L'(\beta)/\sqrt{n}$, which is $\mathcal{O}(1)$.

Now consider the second derivative,

$$L''(\beta) = \sum_{j=1}^{n} \frac{\partial^2}{\partial\beta^2} \log p_j(\mathbf{Y}_j|\beta), \quad (8)$$

which has mean $-n\bar{F}$ where $\bar{F}$ is the average Fisher information on β, defined as $\bar{F} = \frac{1}{n}\sum F$. The total Fisher information is $\mathcal{O}(n)$, so we can define an $\mathcal{O}(1)$ standardised quantity appropriate for using the CLT:

$$\mathcal{Z}^{(2)}(\beta) \equiv \frac{1}{\sqrt{n}}(L''(\beta) + n\bar{F}), \quad (9)$$

which is asymptotically Gaussian with zero mean. Similarly, define a standardised third derivative by

$$\mathcal{Z}^{(3)}(\beta) = \frac{1}{\sqrt{n}}(L'''(\beta) - n\bar{K}), \quad (10)$$

where the $\mathcal{O}(1)$ quantity $\bar{K}$ is given by

$$\bar{K} \equiv \frac{1}{n}\sum_{j=1}^{n} \left\langle \frac{\partial^3}{\partial\beta^3} \log p_j(\mathbf{Y}_j|\beta) \right\rangle. \quad (11)$$

Substituting these definitions into Equation (3), we have

$$\sqrt{n}\mathcal{Z}^{(1)} + \frac{\hat{\delta}}{\sqrt{n}}(-n\bar{F} + \sqrt{n}\mathcal{Z}^{(2)}) + \frac{\hat{\delta}^2}{2n}(n\bar{K} + \sqrt{n}\mathcal{Z}^{(3)}) = 0, \quad (12)$$

or collecting terms of the same order

$$\sqrt{n}(\mathcal{Z}^{(1)} - \bar{F}\hat{\delta}) + (\hat{\delta}\mathcal{Z}^{(2)} + \frac{\hat{\delta}^2}{2}\bar{K}) + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) = 0. \quad (13)$$

Rearranging, we have

$$\hat{\delta} = \frac{\mathcal{Z}^{(1)}}{\bar{F}} + \frac{1}{\bar{F}\sqrt{n}}\left(\hat{\delta}\mathcal{Z}^{(2)} + \frac{\hat{\delta}^2\bar{K}}{2}\right) + \mathcal{O}\left(\frac{1}{n}\right). \quad (14)$$

To leading order, we only need the term in brackets to $\mathcal{O}(1)$, so we replace $\hat{\delta}$ with its zero-order value $\hat{\delta} = \mathcal{Z}^{(1)}/\bar{F}$ to get

$$\hat{\delta} = \frac{\mathcal{Z}^{(1)}}{\bar{F}} + \frac{1}{\bar{F}\sqrt{n}}\left(\frac{\mathcal{Z}^{(1)}\mathcal{Z}^{(2)}}{\bar{F}} + \frac{\mathcal{Z}^{(1)\,2}\bar{K}}{2\bar{F}^2}\right) + \mathcal{O}\left(\frac{1}{n}\right). \quad (15)$$

This is the MLE at next-to-leading order in n. To find its bias, we take the expectation value of $\hat{\delta}$. To compute this, we need the results $\langle\mathcal{Z}^{(1)}\rangle = 0$, $\langle\mathcal{Z}^{(1)\,2}\rangle = \bar{F}$, as well as the definition

$$\bar{J} \equiv \langle\mathcal{Z}^{(1)}\mathcal{Z}^{(2)}\rangle = \frac{1}{n}\sum_{i,j}\left\langle \frac{\partial \log p_i}{\partial\beta}\frac{\partial^2 \log p_j}{\partial\beta^2}\right\rangle. \quad (16)$$

Note that only the diagonal components contribute to the double sum in Equation (16), ensuring that $\bar{J} \sim \mathcal{O}(1)$. Thus we have the bias of the MLE

$$b \equiv \langle\hat{\beta} - \beta\rangle = \frac{1}{2n\bar{F}^2}(\bar{K} + 2\bar{J}) + \mathcal{O}\left(\frac{1}{n^{3/2}}\right), \quad (17)$$

which is the standard result from Bartlett (1953) and subsequent works. It may be generalised to the multi-parameter case to recover the result of Cox & Snell (1968). For Gaussian noise it may be shown that this is actually correct to $\mathcal{O}(n^{-2})$. We denote the leading order part of the bias as $b^{(1)}$.

It is clear from the form of the terms $\bar{K}$ and $\bar{J}$ that a bias arises from non-Gaussianity in the likelihood, in particular, $\bar{K}$ measures the expected skewness of the likelihood while $\bar{J} = -\bar{K} - \bar{F}'$ measures a combination of skewness and position-dependent curvature of the log-likelihood around the maximum in parameter space.

To link the above expression with previous work on noise bias, we note that the 1σ uncertainty on β is proportional to $F^{-1/2} \sim \mathcal{O}(n^{-1/2})$, whereas the leading order bias is $\mathcal{O}(n^{-1})$, making the ratio of the noise bias to the standard deviation generically smaller than unity. However, if the MLE is subsequently averaged over a population of galaxies (as in the case of ellipticities and shear estimates), the error on that average will eventually become smaller than the noise bias. Either the bias must be subtracted, or the posteriors on β from each galaxy must be correctly propagated to ensure the $b/\sigma \sim \mathcal{O}(n^{-1/2})$ behaviour is recovered. We will explore the former technique in this work, noting only that the latter is at the heart of the Bayesian approach to shear estimation.

In order to calculate the bias, we need to evaluate the expectation values in Equation (17) at the true parameter value. However, the difference between the MLE and the true value is $\mathcal{O}(n^{-1/2})$, so to the accuracy of Equation (17) we can evaluate the right-hand-side with the MLE as a proxy for the true parameter value. Sensitivity to this choice indicates the significance of higher-order terms in the expansion of the ML equation and higher-order terms in n.

2.1 Pixel noise and galaxy model assumptions

The noise bias expression Equation (17) does not assume any particular functional form for the likelihood. In the case of shape estimation from a galaxy image, the major sources of pixel noise are expected to be Poisson noise from discrete photon counts, a Gaussian background from unresolved sources, and a small Gaussian read-out noise contribution from the CCD (Mandelbaum et al. 2014). For the remainder of this work, we make the simplifying assumption that the noise is Gaussian and uncorrelated between pixels with variance σ_n^2. The log-likelihood is then

$$L(\beta) = -\frac{1}{2\sigma_n^2}\sum_i [\hat{I}_i - I_i^M(\beta)]^2 + \text{contant}, \quad (18)$$

where $\hat{I}_i$ is the measured surface brightness in pixel i and $I_i^M(\beta)$ is the model prediction at the point β in parameter space. It is straightforward to show in this case that

$$\bar{K} = -3\bar{J} = \frac{3}{n\sigma_n^2}\sum_i \frac{\partial I_i^M}{\partial\beta}\frac{\partial^2 I_i^M}{\partial\beta^2}, \quad (19)$$

where the derivatives are evaluated at the true parameter value. This expression makes it clear that the noise bias arises due to the non-linear mapping between the parameter of interest and the data. A useful expression for the $\mathcal{Z}$-matrices introduced earlier in the derivation of the MLE is, for Gaussian noise,

$$\mathcal{Z}^{(n)}(\beta) = \frac{1}{\sigma_n^2\sqrt{n}} \sum_i [\hat{I}_i - I_i^M(\beta)] \frac{\partial^{(n)} I_i^M}{\partial \beta^{(n)}}. \quad (20)$$

The $\mathcal{Z}$-matrices thus represent weighted linear combinations of the data and are thus Gaussian random variables when the noise is Gaussian.

For the galaxy model, we assume a two-dimensional Gaussian surface-brightness profile with total flux S_T, and covariance matrix

$$\mathbf{Q} = \frac{r^2}{2} \begin{pmatrix} 1+e_1 & e_2 \\ e_2 & 1-e_1 \end{pmatrix}. \quad (21)$$

This galaxy model is not particularly realistic, but will prove sufficient to demonstrate the fundamentals of noise bias and the MLE. Similarly, we do not consider the effects of the PSF on the biases, but note that such effects could easily be incorporated if a model for the PSF were provided. For example, a Gaussian PSF model would introduce an extra additive covariance matrix into Equation (21).

The advantage of analytic functional forms for the likelihood and galaxy model is that the derivatives required to compute the bias in Equation (17) can be done exactly, without the need for finite-differencing, although numerical derivatives could be used in this formalism if required.

2.2 Ellipticity biases from noisy images

To gain further insight into noise bias in the MLE, we pixellize the surface-brightness profile on a coarse grid of 10×10 pixels. The convolution of the surface-brightness profile with the pixel window function is done on an upsampled grid using Fast Fourier Transforms, then downsampled to the observed resolution. The upsampling resolution is adaptively set depending on the typical curvature radius of the input image profile. We have checked that our results are insensitive to changes in the upsampling resolution[1]. The coarse grid allows for rapid calculation of noise biases and likelihood derivatives.

We fix the galaxy scale-length r to 1.2 pixels. A typical pixel size for planned space-based surveys such as *Euclid* is 0.1 arcseconds, corresponding to a galaxy scale-length of 0.12 arcseconds. This is roughly 1σ smaller than the peak of the intrinsic distribution of scale-lengths used by the CFHTLenS survey (Miller et al. 2013) and the 0.3 arcseconds of a typical *Euclid* source (Cropper et al. 2012). Given the steep slope of the mass function for the typical galaxies used in weak lensing surveys, we expect most of the sources in a flux-limited survey to be close to the resolution limit (PSF size) of the telescope, so our assumed value of r does not represent an extreme choice. However it should be borne in mind that this conservative choice implies that our bias results are slightly pessimistic. Combined with our coarse pixellization, this choice of r has the advantage that biases induced by the finite postage-stamp size are avoided. These extra biases, which do not form part of the current study, have recently been speculated as potentially problematic for ML-based methods (GREAT3; Mandelbaum et al. 2015).

We set the peak surface brightness value $I_0 = S_T/2\pi\sqrt{|Q|}$ to unity and then set the pixel noise variance σ_n^2 by the signal-to-noise (S/N), which is defined in the same way as Mandelbaum et al. (2014)

$$S/N = \frac{\sqrt{\sum_i [I_i^M(\beta_0)]^2}}{\sigma_n}, \quad (22)$$

where the reference model $I^M(\beta_0)$ is taken to have an ellipticity of $(e_1, e_2) = (0.3, 0.0)$. Note that we will sometimes use ν to denote S/N in this work. The S/N as defined above clearly scales as $\sqrt{n}$, and it is straightforward to show that the first-order ML bias in Equation (17) scales as ν^{-2}, consistent with previous studies of noise bias (e.g. Hirata et al. 2004).

We compute derivatives by analytically differentiating the Gaussian surface-brightness profile and then pixellizing. For ellipticity biases this requires some care since oscillatory features in the the image domain induced by differentiating require a higher sampling rate in the convolution step.

In the top panel of Figure 1 we plot the bias on e_1 assuming $e_2 = 0$, as a function of the true value of e_1. Given the S/N scaling we choose to plot the ν-independent quantity $b^{(1)}(\nu/10)^2$. The bias is computed over the range $|e_1| \leq 0.98$. Higher values of $|e_1|$ require a prohibitively high sampling rate. Even in this very simplified example a few instructive lessons may be learned. The bias is clearly a highly non-linear function of e_1 for $|e_1| \gtrsim 0.1$. We also note that while antisymmetric, the sign of the bias is never guaranteed to be opposite to that of e_1, i.e. noise bias does not necessarily isotropize the image.

At large $|e_1|$ the bias grows very large compared to its typical $\mathcal{O}(10^{-2})$ values at small $|e_1|$. This behaviour is due to the finite pixellization scale. Since we have chosen the grid to be centred on the galaxy image with an even number of pixels on each side, for increasing values of e_1 the central pixels sample the underlying surface-brightness profile at increasingly higher σ away from the image centre. Due to the rapidly decaying tails of the Gaussian galaxy profile, this means that the derivatives of the image profile with respect to e_1 become very large at large $|e_1|$. This is illustrated schematically in Figure 2.

The effects of finite pixellization may be studied by comparing with the exact continuum-limit (i.e. zero pixel-size) result, which agrees reasonably well with our results for $|e_1| \lesssim 0.6$. In the middle panel of Figure 1 we show this regime, which represents the typical 2σ range of observed ellipticities of galaxies in typical weak lensing surveys. Even in this mildly elliptical regime the bias is highly non-linear, a feature which will be important when we try to remove the bias from the MLE. The rough behaviour of the bias with e_1 matches the continuum-limit over this range of ellipticity, but the amplitude is reduced even for circular galaxies.

In the lower panel of Figure 1 we plot the bias on e_2 as a function of the true value of e_2 assuming $e_1 = 0$. Unlike for e_1, the bias on e_2 remains much closer to its continuum-

[1] Note that the convolution of a Gaussian and a top-hat can be done analytically when the image is aligned with the grid, providing a useful check for numerical errors.

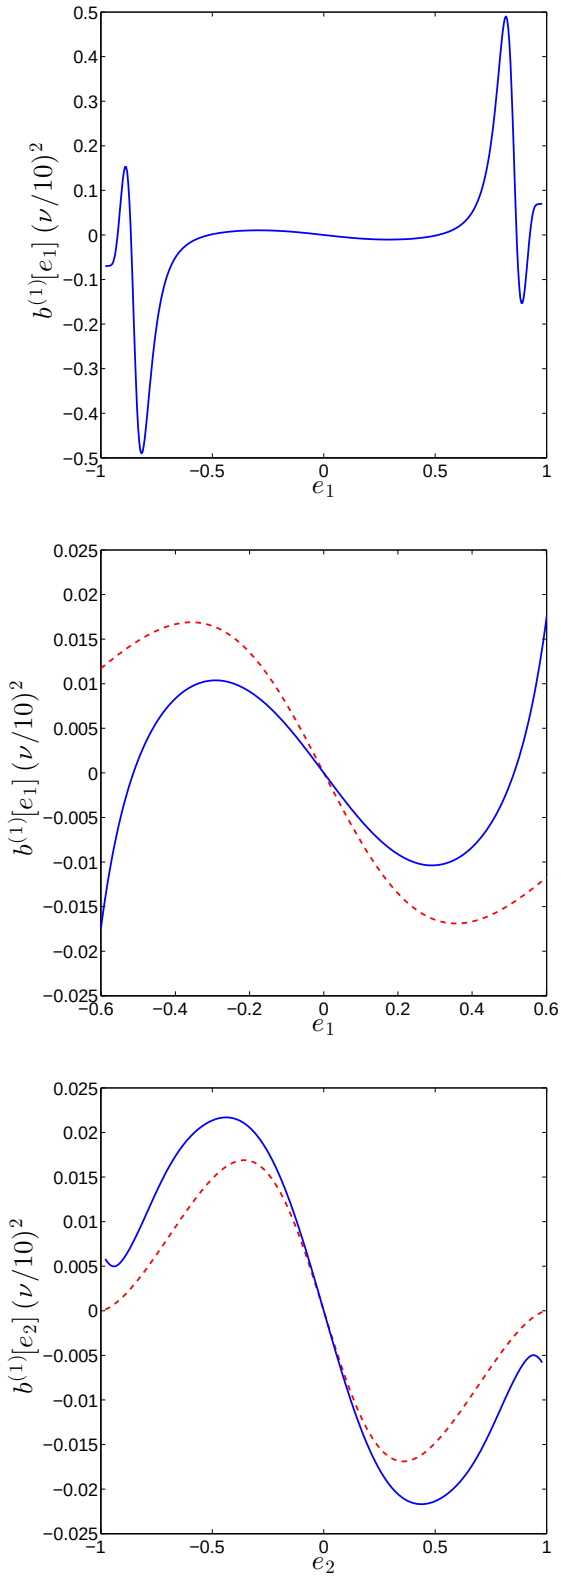

Figure 1. *Upper panel:* Normalized first-order MLE bias on e_1. *Middle panel:* Normalized first-order MLE bias on e_1 (blue solid) and its continuum-limit result (red dashed). *Lower panel:* Normalized first-order MLE bias on e_2 (blue solid) and its continuum-limit result (red dashed).

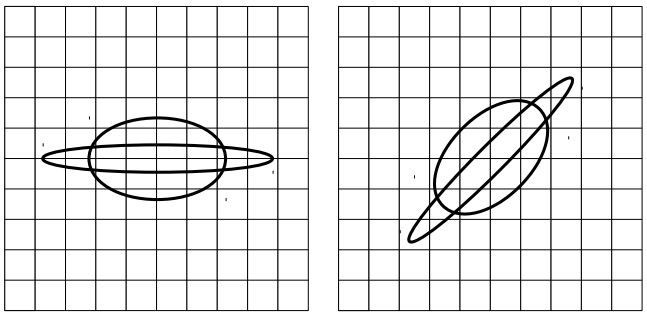

Figure 2. *Left panel:* Effect of increasing e_1 when $e_2 = 0$ and r is held fixed. Solid lines enclose 1σ areas of surface brightness. As e_1 grows, all pixel centres eventually sample the tails of the Gaussian surface brightness profile rather than its core. Small changes to e_1 can thus gives large changes in surface brightness samples, and hence large noise biases. *Right panel:* The equivalent situation for e_2. Now there are always pixel centres sampling the inner regions of the Gaussian. Small changes in e_2 thus give smaller changes to the sampled surface brightnesses, and so the bias is smaller.

limit form out to high values of $|e_2|$. The reason for this is again due to our assumption of a compact Gaussian functional form for the surface brightness profile coupled with the finite pixellization. A galaxy with zero e_1 lies 45° to the horizontal axis, and hence when $|e_2|$ is large there are still pixels sampling the core of the Gaussian along its 'ridge'. Thus the derivatives with respect to ellipticity are never as large as they are for e_1, and the bias is smaller, as illustrated in the right panel of Figure 2.

We emphasise that the pixellization and galaxy size we have chosen for these examples are not extreme choices, and the amplification of noise bias for highly elliptical images could prove a real issue for upcoming weak lensing surveys. One way of mitigating this would be to allow for the centroid to vary and possibly marginalised over, which should go some way to restoring isotropy.

Despite the extreme behaviour of the bias at large $|e|$, we will see that the bias can be still be accurately predicted and removed.

2.3 The bias-corrected MLE and second-order biases

The advantage of the analytic form for the noise bias derived above is that the leading-order bias may be subtracted off without external calibration from simulations. We define the first-order bias-corrected MLE as

$$\hat{\beta}^{(1)} \equiv \hat{\beta} - b^{(1)}(\hat{\beta}), \tag{23}$$

where we have used the MLE as a proxy for the true parameter value in the first-order bias function.

If the MLE is used as a proxy for the true parameter value in this way, the sampling distribution of the resulting *bias-corrected* MLE will acquire some non-Gaussianity due to the non-linear dependence of the ML bias on the data. Thus, it may be thought that this procedure for removing bias might spoil the nice asymptotic properties of the MLE, in particular its efficiency over other point estimators. However it may be shown that removing the $\mathcal{O}(n^{-1})$ bias results in a second-order efficient estimator (Rao 1961). In other words, after correction for the bias, the MLE still has the

optimal variance amongst other point estimators not only at leading order in $\mathcal{O}(n^{-1})$ but at next-to-leading order. This suggests that the bias-corrected MLE is the best one can hope to achieve at this order in $\mathcal{O}(n^{-1})$.

The estimator in Equation (23) is unbiased at $\mathcal{O}(n^{-1})$, or equivalently at $\mathcal{O}(\nu^{-2})$. It does however still receive bias both from higher order terms neglected in the perturbative expansion of the ML equation and from the fact that the MLE differs from the true parameter value. We refer to these second-order biases as the intrinsic and the bias-on-bias terms respectively. For Gaussian noise, both are $\mathcal{O}(\nu^{-4})$. The derivation of the intrinsic term follows straightforwardly from the ML equation using the same techniques as above, so we do not reproduce it here. For the one-parameter case the intrinsic term is, for Gaussian noise,

$$b^{(2)} = \frac{1}{n^2\bar{F}^3}\left[\frac{\bar{K}\bar{P}}{2\bar{F}} + \frac{\bar{K}^3}{24\bar{F}^2} + \frac{\bar{K}\bar{M}}{12\bar{F}} + \frac{\bar{R}}{2} + \frac{\bar{Q}}{40}\right], \quad (24)$$

where

$$\begin{aligned}
\bar{P} &\equiv \frac{1}{n\sigma_n^2}\sum_i \left(\frac{\partial^2 I_i^M}{\partial\beta^2}\right)^2, \\
\bar{R} &\equiv \frac{1}{n\sigma_n^2}\sum_i \frac{\partial^2 I_i^M}{\partial\beta^2}\frac{\partial^3 I_i^M}{\partial\beta^3} \\
\bar{M} &\equiv \frac{1}{n}\langle L^{\mathrm{IV}}\rangle = -\frac{1}{n\sigma_n^2}\sum_i \left[4\frac{\partial I_i^M}{\partial\beta}\frac{\partial^3 I_i^M}{\partial\beta^3} + 3\left(\frac{\partial^2 I_i^M}{\partial\beta^2}\right)^2\right], \\
\bar{Q} &\equiv \frac{1}{n}\langle L^{\mathrm{V}}\rangle = -\frac{5}{n\sigma_n^2}\sum_i \left[\frac{\partial I_i^M}{\partial\beta}\frac{\partial^4 I_i^M}{\partial\beta^4} + 2\frac{\partial^2 I_i^M}{\partial\beta^2}\frac{\partial^3 I_i^M}{\partial\beta^3}\right].
\end{aligned} \quad (25)$$

It is difficult to gain a clear intuition for the various terms in the second-order intrinsic bias, but qualitatively we can say that while the first-order bias arises from skewness in the likelihood, the second-order bias arises from its kurtosis and higher-order non-Gaussian moments.

The bias-on-bias term arising from use of the MLE instead of the true parameter value in $b^{(1)}$ is more straightforward to derive, and can be found either by Taylor expanding $b^{(1)}(\hat{\beta})$ around the true parameter value or by considering $b^{(1)}$ as a parameter and finding the first-order MLE bias using Equation (17). We find the result to be

$$b^{(1)}[b^{(1)}] = \frac{1}{n^2\bar{F}^3}\left[-\frac{13\bar{K}\bar{P}}{48\bar{F}} + \frac{7\bar{K}^3}{27\bar{F}^2} + \frac{13\bar{K}\bar{M}}{48\bar{F}} - \frac{\bar{R}}{4} + \frac{\bar{Q}}{20}\right]. \quad (26)$$

In the left-hand panel of Figure 3 we plot the second-order biases on e_1 as a function of the true value. As in the case of the first-order bias, both second-order terms become large for highly elliptical galaxies. Furthermore, the two contributions are never guaranteed to be of the same sign, meaning that for different values of e_1 the total second-order bias, $b^{(2)} - b^{(1)}[b^{(1)}]$, may be enhanced or reduced through cancellations. For $|e_1| \lesssim 0.75$ however the terms are of the same sign and the bias is reduced.

The second-order bias is much more sensitive to finite-pixellization effects than its first-order counterparts. We find that both second-order terms match their continuum-limit forms well for $|e_1| \lesssim 0.2$, but deviate for more elliptical images. In the middle panel of Figure 3 we zoom in on this region, which demonstrates that for typical galaxy ellipticities the bias-on-bias term is larger than the intrinsic second-order bias. The increased sensitivity to pixellization of second-order terms is due to the higher-order derivatives they require for calculation, coupled with the steepness of the Gaussian surface brightness profile at large $|e_1|$.

In the right-hand panel of Figure 3 we show the second-order biases on e_2 as a function of the true parameter value. As in the case of the first-order bias, the second-order bias on e_2 is less sensitive to the pixellization and matches its continuum-limit value much more closely even for high $|e_2|$. In contrast to the e_1 bias, both second-order terms for e_2 are of the same sign, allowing for cancellation in the total bias on $\hat{\beta}^{(1)}$. Similarly to e_1, the bias-on-bias tends to be larger than the intrinsic term.

With the second-order bias now calculable, we can define a second-order bias-corrected MLE by

$$\hat{\beta}^{(2)} \equiv \hat{\beta} - b^{(1)}(\hat{\beta}) - b^{(2)}(\hat{\beta}) + b^{(1)}[b^{(1)}](\hat{\beta}). \quad (27)$$

This estimator now has a leading-order bias from $\mathcal{O}(\nu^{-6})$ terms, and since the MLE is both third and fourth-order efficient (Kano 1996) we might speculate that $\hat{\beta}^{(2)}$ is still optimal over other second-order bias-corrected point estimators in the sense of having a smaller mean-squared error.

Aside from defining higher-order bias-corrected estimators, the advantage of being able to calculate the second-order bias is that it allows us to approximate the minimum S/N for which the perturbative bias predictions are expected to be valid. We find this by setting the intrinsic second-order bias equal to the first-order bias for each true ellipticity value and solving for the S/N, denoted ν_{min}. In the left-hand panel of Figure 4 we plot this quantity, as well as the S/N at which the second-order bias is 10% of the first-order bias. The spikes in this plot arise from points where the first-order bias goes to zero, at which point higher-order terms necessarily become significant. It is interesting to note that this can happen even for mildly elliptical galaxies with $e_1 \approx 0.5$. Zooming in on small $|e_1|$ in the middle panel of Figure 4, we see that the minimum useable S/N is quite flat in e_1, with second-order terms becoming significant at $S/N \approx 10$. The situation is even better for e_2, with the right-hand panel of Figure 4 suggesting that $S/N \approx 10$ is a reasonable approximation to the minimum S/N for which our perturbative bias formalism will be accurate.

3 ELLIPTICITY BIASES FROM SIMULATED IMAGES

Now we have the machinery to predict the MLE bias directly from the likelihood, we can test the performance of the bias-corrected estimates on simulated noisy galaxy images. For this purpose we generate a large number of pixellized galaxy images with random Gaussian noise, and a true ellipticity having either $(e_1, e_2) = (0.3, 0.0)$ when using e_1 as a parameter or $(e_1, e_2) = (0.0, 0.3)$ when using e_2. At this stage we stick to one-dimensional parameter inferences. For each noisy image we find the MLE by gridding up the likelihood, and then compute the first and second-order biases with the MLE as a proxy for the true parameter value. This allows us to construct first and second-order bias corrected estimators. We have checked that our results are insensitive to changes in the likelihood grid-spacing when finding the maximum.

In the left-hand panel of Figure 5 we show the resultant

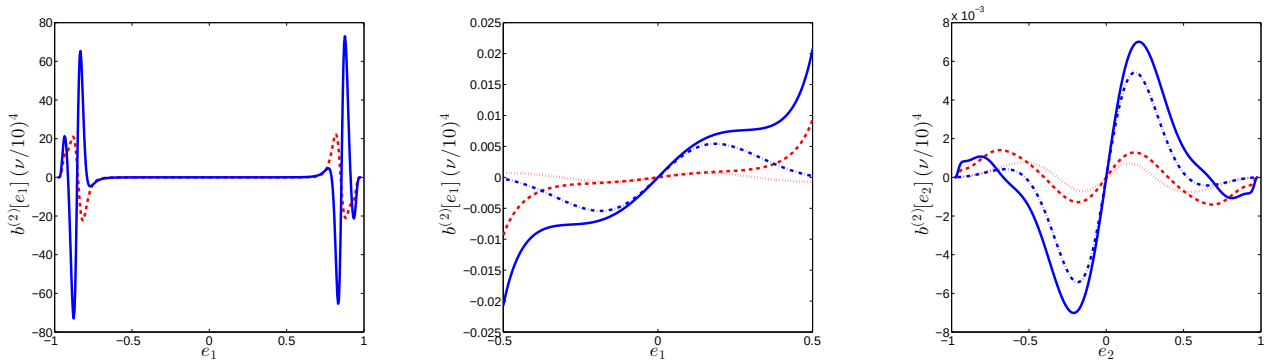

Figure 3. *Left panel:* Normalized second-order MLE biases on e_1, showing the bias-on-bias term (blue solid) and the intrinsic term (red dashed). *Middle panel:* Normalized second-order MLE biases on e_1, showing the bias-on-bias term (blue solid), its continuum-limit result (blue dot-dashed), the intrinsic term (red dashed) and its continuum-limit result (red dotted). *Right panel:* Same as middle panel for e_2.

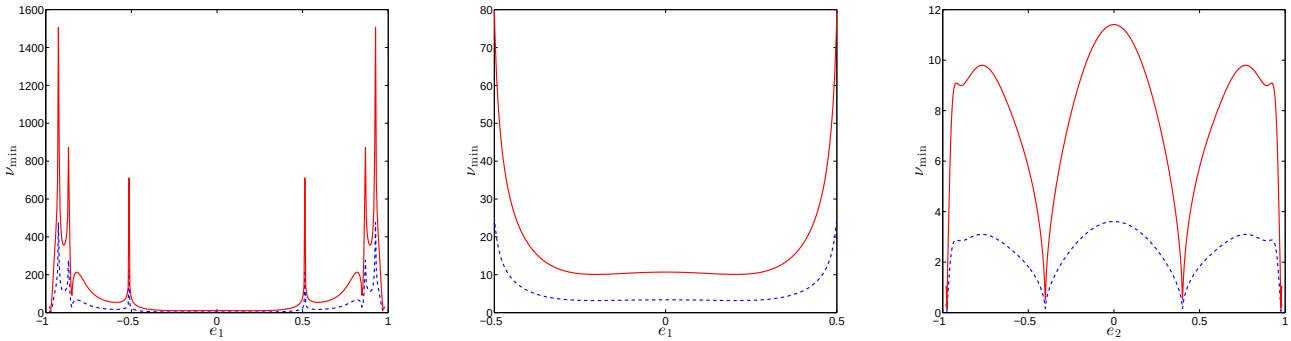

Figure 4. *Left panel:* Value of the S/N at which the intrinsic second-order bias on e_1 equals the first-order MLE bias as a function of the true parameter value (blue dashed) and the S/N at which it is 10% of the first-order bias (red solid). *Middle panel:* Same as left panel for small $|e_1|$. *Right panel:* Same as left panel for e_2.

biases on e_1 for the uncorrected, first-order corrected, and second-order corrected estimators for a range of S/N values. The error bars are calculated from the empirical variance of the MLE for each noise realization at the given S/N. We plot the magnitude of the dimensionless quantity b/e_0 where e_0 is the true ellipticity. Values of unity on the horizontal axis thus represent order unity changes to the measured ellipticity from noise bias.

From this figure we see that a first-order bias correction can reduce the bias on the MLE by several orders of magnitude at high S/N, and helps to reduce bias all the way down to $S/N = 15$. At lower S/N, higher-order terms become significant and a first-order correction is not sufficient. The second-order correction does even better, clearly reducing the bias further down to $S/N = 12$. For noisier images a second-order correction actually increases bias, due to the importance of higher-order terms in the expansion of the ML equation, as expected from Figure 4.

To gain further insight into the performance of the estimators, in the left-hand panel of Figure 6 we plot the empirically measured biases from the simulations against their theoretical expectations. The dashed lines in this figure are the predictions from Section 2 using the true parameter value, and the solid lines are the empirical estimates of these biases using the MLE as a proxy for truth. Thus the dashed lines scale precisely as ν^{-2} and ν^{-4}. At high S/N the MLE exhibits small fluctuations about the true parameter value, and hence the measured biases closely match the predictions. As we increase the noise variance, the MLE starts to depart more frequently from the truth in each noise realization. The difference between the mean of the measured first-order biases and the true values is given, to leading order, by the bias-on-bias term introduced above. From Figures 1 and 3 we see that at $e_1 = 0.3$ this term is positive, whereas the first-order bias is negative. Thus, as the S/N is decreased, the mean first-order bias measured from the simulations becomes less negative than the true value, and so the red points in Figure 6 underpredict the true values give by the red dashed line. Around $S/N = 10$, the second-order bias-on-bias (blue dashed line) is comparable to the first-order bias, eventually exceeding it and causing a sign-change in the red points at $S/N = 5$. This explains the 'kink' in the red points at $S/N = 10$ in the left-hand panel of Figure 6.

Due to the highly non-linear dependence of the biases on e_1 (Figures 1 to 3), the distribution of measured biases becomes highly non-Gaussian, with a mean that departs significantly from the analytic prediction. For very low S/N, the MLE can be in the region of e_1 parameter space where the biases become very large. These large values of the bias

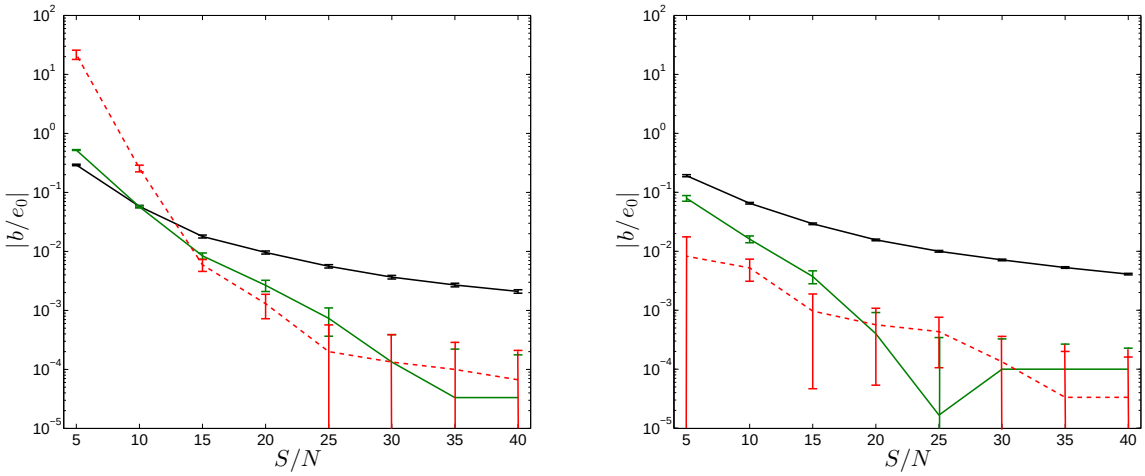

Figure 5. *Left panel:* Magnitude of measured bias on the e_1 MLE compared to the true value (black solid), the first-order bias-corrected MLE (green solid) and the second-order bias-corrected MLE (red dashed). *Right panel:* Same as left panel for e_2.

can drag up the mean value significantly, and this non-Gaussianity renders the error bars on these plots unreliable at low S/N.

Another way of understanding this behaviour is by recalling that noise bias arises from non-Gaussianities in the likelihood. The extremely non-linear relationship between the biases and the ellipticity induces very large bias-on-bias terms, which explains the behaviour at low S/N.

In the right-hand panel of Figure 5 we plot the measured MLE biases on e_2 as a function of S/N, for a fiducial model of $(e_1, e_2) = (0.0, 0.3)$. At high S/N, independent noise realizations give rise to MLEs that exhibit only small fluctuations around the true parameter value. In this regime, the bias is given roughly by the value of curve in the bottom panel of Figure 1 at the true value $e_2 = 0.3$. This is roughly twice the corresponding bias in e_1 at the same S/N (middle panel of Figure 1), so the measured bias at high S/N in the right-hand panel of Figure 5 is roughly twice as large as the corresponding points in the left-hand panel of Figure 5. At low S/N this behaviour is reversed, the bias in e_1 now being dominated by the large-$|e_1|$ behaviour seen in the upper panel panel of Figure 1, such that the bias in e_2 is now less than that of e_1.

Furthermore, since neither the first or second-order e_2 biases grow very large (Figures 1 and 3), the perturbative expansion of the ML equation is much more accurate than for e_1 at low S/N, and hence a first-order correction always improves the estimate, reducing the bias by an order of magnitude or so down to $S/N = 15$ and by a factor of a few down to $S/N = 5$. A second-order correction does even better, giving more than an order-of-magnitude improvement even at $S/N = 5$, with third-order terms expected to be subdominant even here according to Figure 4. This point is backed up in the right-hand panel of Figure 6, which clearly demonstrates that the measured biases track their predictions more closely than their e_1 counterparts.

The main conclusion of this section is that the MLE noise bias can be reduced by several orders of magnitude for $S/N \gtrsim 30$ with our method, and is reduced for all $S/N \gtrsim 15$. The improvement is better for e_2 than e_1, due to the more linear relationship of the bias to the ellipticity, and the smaller intrinsic biases caused by the coarse pixellization. Second-order corrections offer further improvement, accounting for much of the bias non-linearity, although these fail too below $S/N \approx 12$ for e_1. For e_2 second-order corrections can reduce noise bias by almost an order of magnitude even at $S/N \approx 5$. We emphasise that these biases have been reduced without any need for external calibration.

4 MAXIMUM A POSTERIORI ESTIMATES

We have seen that the estimation of the MLE bias from a coarsely pixellized noisy galaxy image is made difficult by the extreme behaviour at large values of the ellipticity. However, we know that the galaxies in weak lensing surveys do not exhibit these strong ellipticities, with the r.m.s. ellipticity being roughly 0.3. It may seem strange then that we should allow values of the MLE to take on such extreme values given our prior knowledge of the intrinsic source distribution. It may then be hoped that introducing a prior on intrinsic ellipticity will regularize the bias-estimation procedure as well as more faithfully representing our strong prior expectations on the true noise-free galaxy shape.

With the introduction of a prior, the quantity we now seek to maximize is the *posterior* rather than the likelihood, and the MLE now becomes the MAP. In this section we investigate the performance of the MAP in the problem of shape estimation from noise images.

4.1 Bias in the MAP

Introducing a prior to regularize bias-estimation comes at the cost of introducing an extra source of bias into the MLE (now the MAP). This can be straightforwardly accounted for using the formalism of Section 2. Defining the logarithm of the prior probability on β as L_p, the log-likelihood is now modified as $L \to L + L_p$. We find the leading-order correction to the first-order noise bias (valid for any noise distribution)

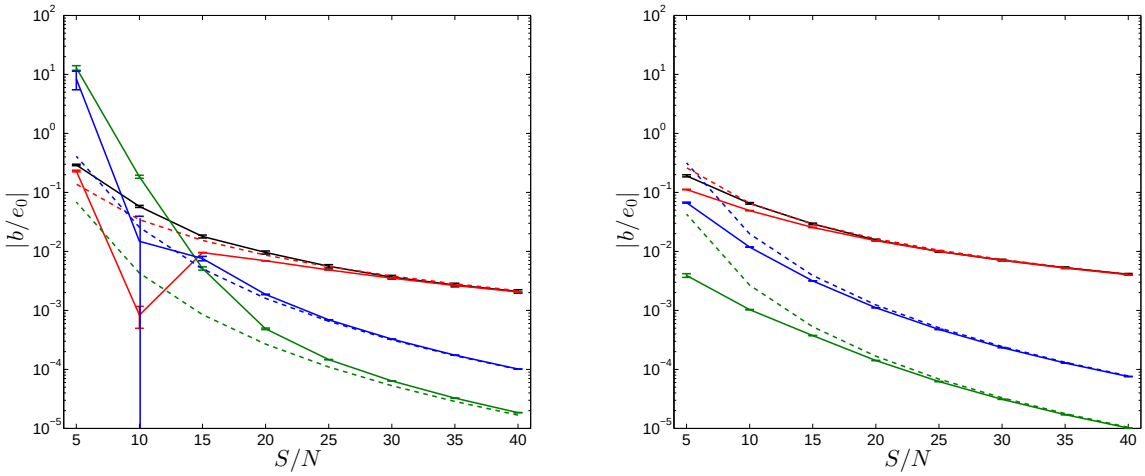

Figure 6. *Left panel:* Magnitude of MLE bias compared to true value of e_1 (black solid), mean value of measured first-order bias (red solid), true first-order bias prediction (red dashed), mean value of measured intrinsic second-order bias (green solid), true intrinsic second-order bias (green dashed), mean value of measured bias-on-bias (blue solid), true bias-on-bias (blue dashed). *Right panel:* Same as left panel for e_2.

to be

$$b_p^{(1)} = \frac{L'_p}{n\bar{I}}. \tag{28}$$

The second-order intrinsic bias receives an extra contribution given (for Gaussian noise) by

$$n^2\bar{I}^3 b_p^{(2)} = L'_p\bar{P} + \frac{L'_p\bar{K}^2}{2\bar{I}} + \frac{5L''_p\bar{K}}{6} + \frac{(L'_p)^2\bar{K}}{2} + \frac{L'_p\bar{M}}{4} + L'_pL''_p\bar{I} + \frac{L'''_p\bar{I}}{2}, \tag{29}$$

while the bias-on-bias second-order term receives the contribution

$$n^2\bar{I}^3 b^{(1)}[b^{(1)}]_p = -\frac{3L'_p\bar{P}}{8} + \frac{7L'_p\bar{K}^2}{9\bar{I}} + \frac{5L''_p\bar{K}}{6} + \frac{2(L'_p)^2\bar{K}}{3} + \frac{3L'_p\bar{M}}{8} + L'_pL''_p\bar{I} + \frac{L'''_p\bar{I}}{2}. \tag{30}$$

We choose for the prior the function

$$p(|e|) \equiv \exp(L_p) \propto (1-|e|^2)\exp\left(-|e|^2/2\sigma_e^2\right). \tag{31}$$

In this section we fix $\sigma_e = 0.3$, deferring discussion of the sensitivity to the prior until Section 6. This form of the prior satisfies the requirement of zero weight at the boundary $|e|=1$.

In the left-hand panel of Figure 7 we plot the total (prior + noise) bias on the MAP for e_1 as a function of the true value, which should be compared to Figure 1. In the middle panel of Figure 7 we zoom on the observationally relevant range of small $|e_1|$, which should be compared to Figure 1. We see from these plots that the prior adds a significant bias to the MLE, and actually diverges at the boundary $|e_1|=1$. This is partly due to the pre-factor multiplying the Gaussian part of Equation (31), whose log-derivative diverges, and partly due to the Fisher information in the denominator of Equation (28) which suffers from the same finite pixellization issues as the MLE. However, since the prior naturally down-weights these extreme ellipticity regions when finding the most likely parameter value, this divergence will not prove problematic.

In the right-hand panel of Figure 7 we plot the total first-order bias on e_2. The main effect of the prior here is to increase the bias on all scales, but the shape is rather similar to the prior-free case (Figure 1). This suggests that the finite-pixellization effect from $\bar{F}$ is more important in determining the high-$|e|$ behaviour of the prior bias than the polynomial pre-factor in Equation (28), which gives us some confidence that our results will be broadly insensitive to the precise form of the prior.

In the left-hand panel of Figure 8 we plot the second-order biases induced by the prior on e_1, with the middle panel of Figure 8 zooming in on small $|e_1|$. These plots show similar features to their prior-free counterparts, in particular the dramatic increase in bias towards the boundaries of the parameter space. However, a key difference is that the second-order terms are now of much more similar magnitude and have the same sign for all e_1. This may be understood by examination of the individual terms in the second-order bias expressions. The dominant terms in both expressions for our choice of prior and galaxy model turn out to be the $L'_pL''_p\bar{I}$ and $5L'''_p\bar{K}/6$ terms, which appear in both second-order bias expressions and hence cancel perfectly. The other important terms have the same sign and similar magnitude, leading to broadly similar values at all e_1 and hence a total bias that is much smaller than either individual term. Similar behaviour is observed in the right-hand panel of Figure 8, where we plot the second-order biases on e_2.

In the left-hand panel of Figure 9 we plot the minimum S/N that can be used for the e_1-MAP with our perturbative bias scheme. The spikes are generally lower in magnitude to their MLE counterpart (Figure 4), as the prior bias prevents the first-order term from going through zero. Zooming in on small e_1 in the middle panel of Figure 9, we see that the extra bias induced by the prior has increased $\nu_{\rm min}$. At the r.m.s. ellipticity of $e_1 = 0.3$ we have $\nu_{\rm min} \approx 20$. This number

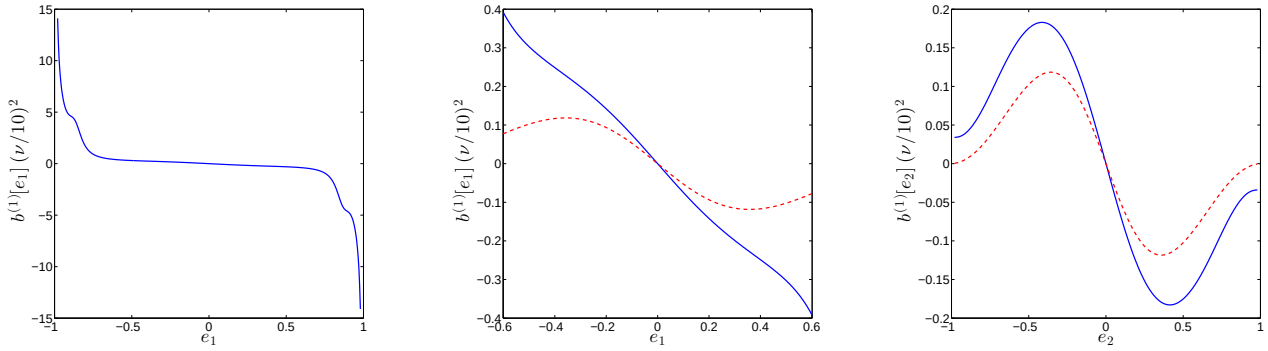

Figure 7. *Left panel:* Normalized first-order MAP bias on e_1. *Middle panel:* Normalized first-order MAP bias on e_1 (blue solid) and its continuum-limit result (red dashed). *Right panel:* Normalized first-order MAP bias on e_2 (blue solid) and its continuum-limit result (red dashed).

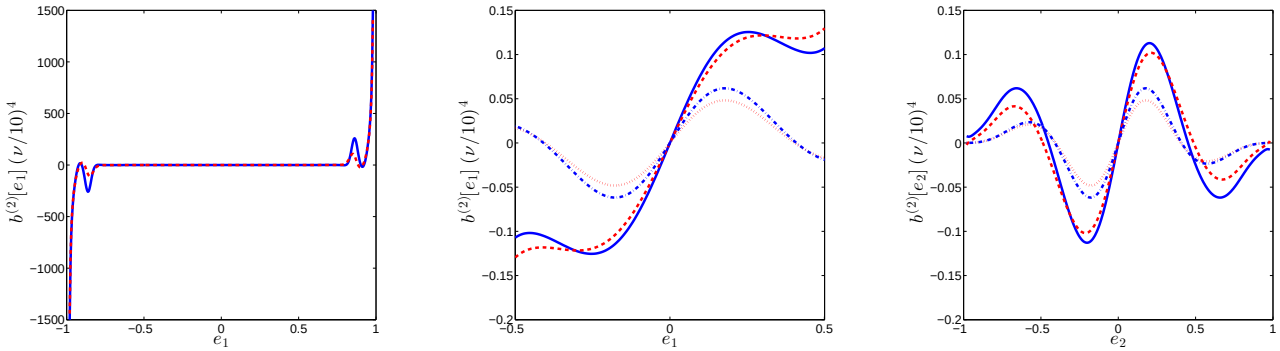

Figure 8. *Left panel:* Normalized second-order MAP biases on e_1, showing the bias-on-bias term (blue solid) and the intrinsic term (red dashed). *Middle panel:* Normalized second-order MAP biases on e_1, showing the bias-on-bias term (blue solid), its continuum-limit result (blue dot-dashed), the intrinsic term (red dashed) and its continuum-limit result (red dotted). *Right panel:* Same as middle panel for e_2.

is rather similar in the case of the e_2-MAP, for which we plot $\nu_{\rm min}$ in the right-hand panel of Figure 9.

4.2 Testing the MAP on image simulations

With predictions for the extra prior-induced bias now at our disposal, we may proceed with repeating the tests of Section 3 with the MAP as our shape estimator instead of the MLE.

In the left-hand panel of Figure 10 we plot the measured total bias (noise + prior) on the MAP from image simulations. This should be compared to Figure 5 for the case of the MLE. The black points in this figure show the bias in the MAP as measured from the image simulations, which is now uniformly larger than that of the MLE due to additional bias brought by the prior. However, both the first-order and second-order bias-corrected estimators succeed in removing most of this extra bias, with the resultant residual biases smaller than their MLE counterparts. Note that the error bars in this plot are too large to allow for a comparison between the first and second-order estimators.

The major difference between Figure 10 and Figure 5 however is the performance at low S/N. In this regime the MLE suffers from large biases due to the large ellipticities that are favoured in chance realizations of the noise. When a prior is included these regions of parameter space are down-weighted, regularizing the perturbative series and 'linearizing' the bias, such that even at $S/N = 10$ a first-order correction reduces the bias by roughly two orders of magnitude from $\mathcal{O}(10^{-1})$ to $\mathcal{O}(10^{-3})$. The same correction applied to the MLE actually worsens the bias. This demonstrates that although the prior brings an extra source of bias, its regularizing effects on our perturbative scheme vastly outweigh this. Across the entire S/N range in Figure 10, the bias-corrected MAP reduces the bias by at least two orders of magnitude for all S/N.

This argument is backed up by inspection of the individual biases as estimated from the image simulations, which are displayed in the left-hand panel of Figure 11, which should be compared with Figure 6 in the MLE case. The solid lines (measured biases) now track the dashed lines (exact biases) much more closely across the S/N range, indicating that the prior is successfully regularizing the perturbative ML expansion. Note that the green and blue lines lie on top of each other in this plot, which is due to the very similar magnitudes of the second-order biases as discussed above. The total second-order bias is the difference in magnitude of these terms, such that the total second-order bias

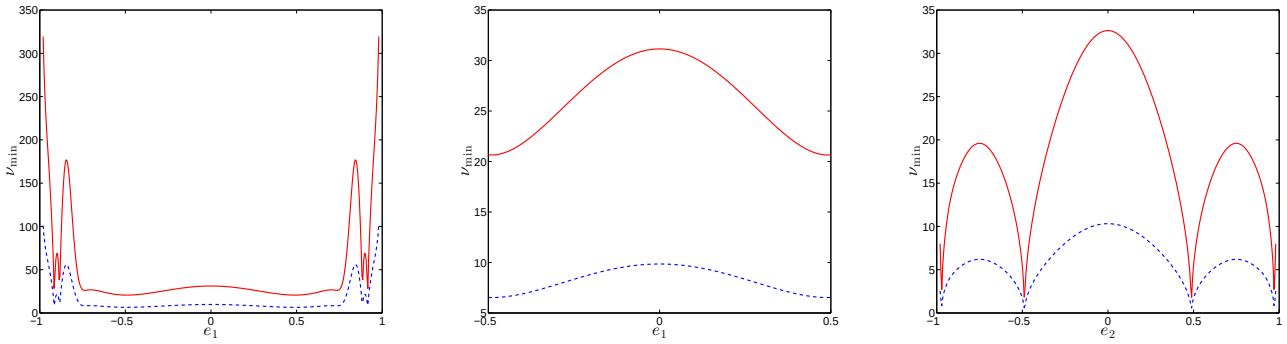

Figure 9. *Left panel:* Value of the S/N at which the intrinsic second-order bias on e_1 equals the first-order MAP bias as a function of the true parameter value (blue dashed) and the S/N at which it is 10% of the first-order bias (red solid). *Middle panel:* Same as left panel for small $|e_1|$. *Right panel:* Same as left panel for e_2.

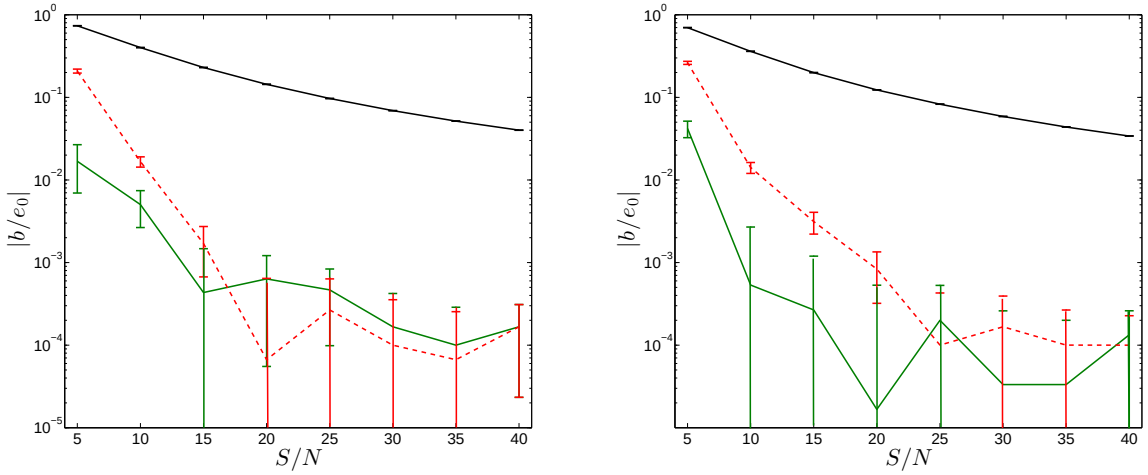

Figure 10. *Left panel:* Magnitude of measured bias on the e_1 MAP compared to the true value (black solid), the first-order bias-corrected MLE (green solid) and the second-order bias-corrected MLE (red dashed). *Right panel:* Same as left panel for e_2.

is much less significant in the case of a prior than without. This is exhibited in the close tracking of the first-order bias (red points) by the total measured bias (black points) in Figure 11, where the lines are almost indistinguishable. This suggests that second-order effects are subdominant in the presence of a prior, due to both the regularization of the ML expansion and the cancellation between intrinsic and bias-on-bias terms. This cancellation is potentially a consequence of our choice of Gaussian noise, and possibly also due to our Gaussian galaxy model assumption, although further investigation of this would require explicit checking against alternative models. However, our results definitively show that the regularizing effect of the prior brings bias measurements into considerably better agreement with their exact expressions.

We have seen that, unlike e_1, the leading-order biases in e_2 do not grow large at large $|e|$. This suggests that introducing a prior will not have as much of an impact on the MLE bias, which is backed up in the right-hand panel of Figure 10 where we plot the measured biases on the e_2 MAP. The improvement in the estimator bias is not as dramatic in the first-order-corrected MAP as it is for the first-order-corrected MLE, although there is still an improvement for all S/N. Furthermore, the second-order-corrected MAP actually performs worse than the first-order-corrected MAP at $S/N \lesssim 15$, in contrast to the MLE. We attribute this to the fact that the ratio of second-order to first-order bias is larger for e_2 with a prior than without (compare Figure 4 and Figure 9). This suggests that the perturbative expansion of the ML equation is likely to break down at a higher S/N with a prior than without.

In the right-hand panel of Figure 11 we plot the individually measured biases on the e_2 MAP. We again see that the bias-on-bias and intrinsic second-order biases are of very similar magnitude, suggesting that the total second-order bias is reduced compared to the naive expectation. Of course, this does not affect the argument above that the a first-order correction is not as effective for the e_2 MAP, which is based on the ML expansion and hence only on the size of the intrinsic second-order bias.

The results of this section strongly suggest that introducing a prior into the MLE can reduce its noise bias by orders of magnitude even at low S/N. Even though the prior itself is a source of additional bias, this too can be easily ac-

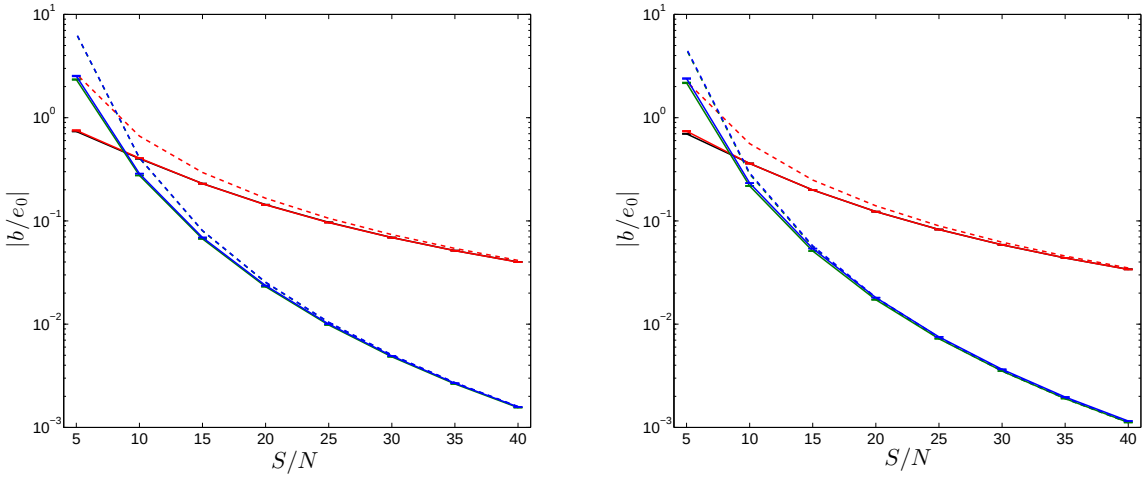

Figure 11. *Left panel:* Magnitude of MAP bias compared to true value of e_1 (black solid), mean value of measured first-order bias (red solid), noise-free first-order bias prediction (red dashed), mean value of measured intrinsic second-order bias (green solid), noise-free intrinsic second-order bias (green dashed), mean value of measured bias-on-bias (blue solid), noise-free bias-on-bias (blue dashed). Black and red points lie on top of each in this plots, as do blue and green points. *Right panel:* Same as left panel for e_2. Black and red points lie on top of each other, as do blue and green points.

counted for in our formalism. The new estimator, the MAP, is more robust to second-order corrections than the MLE and little is gained by the inclusion of such terms, although they do offer a very useful test of the convergence of the ML equation to its asymptotic form. Although our tests have thus far been on simple galaxy models with Gaussian noise, the formalism we have developed could be extended to more complicated likelihood functions. Correlated noise could in principle also be easily included, although sufficiently many pixels would need to sample the bright regions of the galaxy to ensure the CLT converges efficiently. We defer investigation of this issue to a future work, and instead now apply our formalism to biases on the measured shear field itself.

5 ONE-DIMENSIONAL SHEAR BIASES

The previous sections have given us an insight into the various factors determining noise bias in galaxy images. However, the relevant observable for gravitational lensing is not the galaxy ellipticity but the shear that acts upon it. In this section we investigate the effects of noise bias on cosmic shear estimates. In this section we will specialise to one-dimension, to facilitate examination of the propagation of the noise bias from the one-dimensional ellipticity parameter spaces discussed in the previous section. We will also largely neglect second-order biases in this section, as the image simulations have shown that they improve biases only modestly over a first-order correction.

Under our definition of ellipticity, the unlensed complex ellipticity e_s is related to the lensed ellipticity e and the shear g by (Schneider & Seitz 1995)

$$e = \frac{e_s + 2g + g^2 e_s^*}{1 + |g|^2 + 2\mathrm{Re}(g e_s^*)}. \tag{32}$$

Restricting to weak shear $g \ll 1$ and aligned shear and ellipticity such that $e = \mathrm{Re}(e) = e_1$ and $g = \mathrm{Re}(g) = g_1$, we have

$$e \approx e_s + 2g(1 - e_s^2), \tag{33}$$

and hence

$$g = \frac{e - e_s}{2(1 - e_s^2)}. \tag{34}$$

Therefore, given a noisy estimate of the lensed ellipticity $\hat{e}$, and the true unlensed ellipticity, a natural estimator of shear is given by

$$\hat{g} = \frac{\hat{e} - e_s}{2(1 - e_s^2)}. \tag{35}$$

It may be wondered whether one could attempt to estimate g directly from an individual galaxy image and hence make use of the bias-correction schemes developed in the preceding sections. Unfortunately, even with a prior on intrinsic ellipticity and the restriction of $|g| < 1$ (Schneider & Seitz 1995) such an estimate would still have infinite noise, i.e. there would be a direction in (g_1, g_2) space such that the posterior probability was constant. To see this, note that an observation of galaxy shape produces two observables (e_1 and e_2), whereas there are four unknowns (both components of e_s and g). A prior on $|e_s|$ reduces this to three unknowns, and so the problem is still unconstrained. Our perturbative scheme is therefore unlikely to succeed when applied to such a shear estimator, so we specialise here to the more common approach of averaging over galaxies in a coherence patch of the shear.

Since the intrinsic ellipticity is unobservable, we average ellipticity estimates over sources to estimate shear. Equation (34) suggests the following form:

$$\hat{g} = \frac{\langle \hat{e} \rangle_{\mathrm{S}}}{2(1 - \sigma_e^2)}, \tag{36}$$

where σ_e^2 is the variance of the intrinsic ellipticity distribution and we have averaged over sources. We have also assumed that e_s has zero mean. The individual ellipticity

estimates are biased by noise, which can be written schematically as

$$\hat{e} \approx e + b^{(1)}(e, r^2, S_T)X^2 \tag{37}$$

where X is a Gaussian random variable with zero mean and unit variance, and we have neglected higher-order noise bias terms. We have made the dependence of the bias on the other galaxy model parameters explicit in the above expression for reasons which will become clear. To average this expression over sources, we need to substitute g and e_s for e:

$$\langle \hat{e} \rangle_{\rm S} = 2g(1-\sigma_e^2) + \langle b^{(1)}[\mathbf{\Theta}(\mathbf{\Theta}_s, g)] \rangle_{\rm S} \langle X^2 \rangle_{\rm S}, \tag{38}$$

where we have assumed that the noise is uncorrelated between galaxies, and $\mathbf{\Theta}$ represents the lensed galaxy model parameters, with $\mathbf{\Theta}_s$ their unlensed values. Now, since we have assumed that the ellipticity biases are separable in e and X (which is true at each order in signal-to-noise), averaging over sources and averaging over noise are commutative operations. So, the shear estimate averaged over noise is

$$\langle \hat{g} \rangle_n = g + \frac{\langle b^{(1)}(\mathbf{\Theta}_s, g) \rangle_{\rm S}}{2(1-\sigma_e^2)}. \tag{39}$$

Finally, we linearize the bias functions around zero shear:

$$\begin{aligned} b[\mathbf{\Theta}(\mathbf{\Theta}_s, g)] &\approx b[\mathbf{\Theta}(\mathbf{\Theta}_s, 0)] + g \left. \frac{\partial b[\mathbf{\Theta}(\mathbf{\Theta}_s, g)]}{\partial g} \right|_{g=0} \\ &= b(\mathbf{\Theta}_s) + g \left. \frac{\partial b(\mathbf{\Theta})}{\partial \Theta_i} \right|_{\mathbf{\Theta}=\mathbf{\Theta}_s} \left. \frac{\partial \Theta_i}{\partial g} \right|_{g=0}, \end{aligned} \tag{40}$$

where we have implicitly summed over the model parameters Θ_i. Defining the shear bias parameters by $\langle \hat{g} \rangle = (1+m)g + c$ we have

$$\begin{aligned} m &= \frac{1}{2}(1-\sigma_e^2)^{-1} \left\langle \frac{\partial b(\mathbf{\Theta})}{\partial \Theta_i} \frac{\partial \Theta_i}{\partial g} \right\rangle_{\rm S} \\ c &= \frac{1}{2}(1-\sigma_e^2)^{-1} \langle b(\mathbf{\Theta}_s) \rangle_{\rm S} = 0, \end{aligned} \tag{41}$$

where the final equality follows since both first and second order bias functions are antisymmetric in the unlensed model parameters. This expression can be evaluated using the known relationship of the lensed model parameters to the shear, and by differentiating the biases with respect to the model parameters then averaging the product over one component of the ellipticity. For the intrinsic ellipticity distribution we take the same function used as the ellipticity prior, Equation (31), although these two functions need not be the same in practice.

Equation (41) demonstrates that it is the derivatives of the bias functions with respect to model parameters that determine the shear biases, rather than just the values of the functions themselves. It also shows that noise bias only contributes to the multiplicative bias and not the additive bias, which is a consequence of isotropy, i.e. the fact that the intrinsic distribution depends only on the magnitude of ellipticity and not its direction.

In the upper panel of Figure 12 we plot the true 1D multiplicative first and second-order shear biases assuming shear has been measured from a population with intrinsic ellipticities given by Equation (31) using only the e_1 component of measured ellipticity, assuming that $e_2 = 0$. This plot shows that the intrinsic second-order bias is subdominant to the first-order term only for $S/N \gtrsim 40$. For noisier images, the large e_1 ellipticity biases contribute to the average in Equation (41), even after downweighting by the intrinsic distribution of Equation (31). At these low S/N values the perturbative expansion of the ML equation will start to break down, and so we predict that these bias predictions will not accurately describe the bias measured from image simulations.

The shear multiplicative biases m and c are useful quantities for us to study because weak lensing surveys usually couch the performance of their shear estimators in terms of these numbers. Planned space-based surveys such as the *Euclid* mission will require the total multiplicative shear bias to be no greater than 2×10^{-3} in order to gain percent-level accuracy on the dark energy equation of state (Laureijs et al. 2011; Massey et al. 2013). From Figure 12 we see that in the regime where our perturbative bias predictions are expected to be accurate ($S/N \gtrsim 40$), the multiplicative bias from noise is $\mathcal{O}(10^{-4})$, so we will have to investigate through image simulations whether our bias-correction schemes might be useful for our fiducial space-based survey at lower S/N.

In the lower panel of Figure 12 we plot the first and second-order bias predictions for the MAP e_1 estimator, averaged over sources. Similarly to the ellipticity biases, we see that the biases are all uniformly larger across the S/N range, due to the extra bias induced by the prior. We also see that the individual second-order terms are of similar magnitude, and have the same sign, a consequence of the similar forms of the MAP ellipticity biases. In contrast to the upper panel of Figure 12, the second-order terms are subdominant to the first-order bias for almost all $S/N \gtrsim 10$, suggesting that the MAP-based shear estimator might be more amenable to our perturbative bias correction than its MLE counterpart.

Using our image simulations, we can form the shear estimate Equation (36), and measure the resulting bias. We apply a uniform weak shear of $g = 0.01$ and simulate galaxies with $e_2 = 0.0$ and e_1 sampled from the intrinsic distribution Equation (31). It is straightforward to show that corrections to m from higher-order terms in g enter at $\mathcal{O}(g^2)$, which can also be seen through symmetry considerations (Madhavacheril et al. 2015).

After finding the MLE (or MAP) for each galaxy, we may apply a first-order (or second-order) correction to that estimate and then proceed to average this 'per-galaxy corrected' estimator. The resulting residual bias for the MLE is shown as the red points in Figure 13, with the blue points showing the uncorrected bias. Alternatively, if we know the intrinsic ellipticity distribution, we may perform a 'global correction' by averaging the uncorrected ellipticity estimates and then subtracting the predicted first-order ellipticity bias averaged over this distribution, which is a known deterministic function, Equation (41). Schematically, the globally corrected estimator is

$$\hat{g}_{\rm GL} \propto \langle \hat{e} \rangle_S - \langle b(e_s) \rangle_S, \tag{42}$$

whilst the per-galaxy corrected estimator is

$$\hat{g}_{\rm PG} \propto \langle \hat{e} - b(\hat{e}) \rangle_S. \tag{43}$$

At leading order in S/N these procedures are equivalent but at second-order they are not. Subtracting biases on each galaxy leaves two second-order residuals (which we have referred to as the intrinsic and bias-on-bias terms), whilst subtracting a global correction leaves us only with one second-order residual (the intrinsic term). The difference between

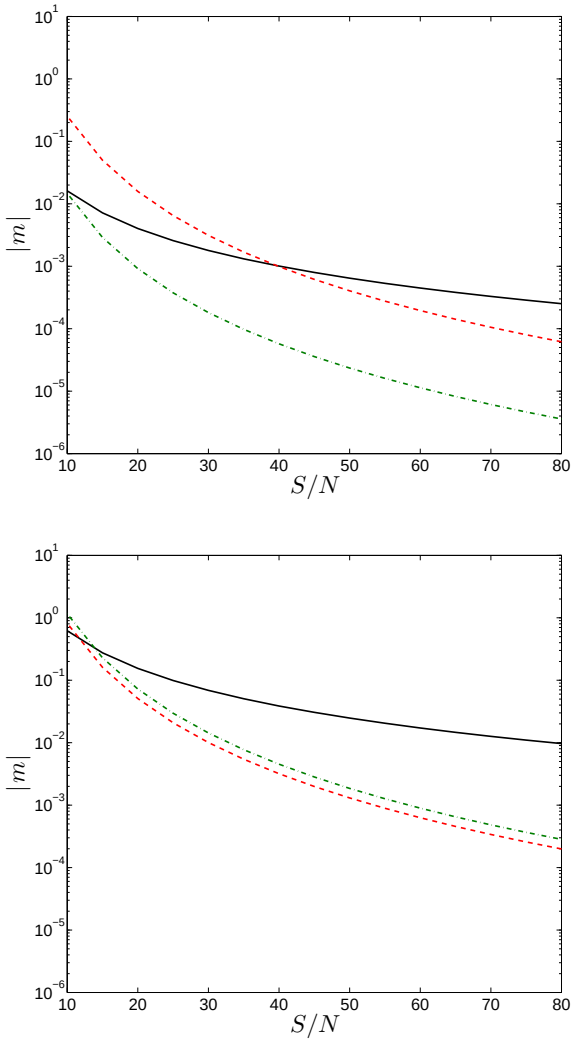

Figure 12. *Upper panel:* Theoretical absolute 1D shear multiplicative MLE biases from e_1, derived from first-order ellipticity bias (black solid), the second order bias-on-bias term (green dot-dashed), and the intrinsic second-order bias (red dashed). Note that the first-order bias is negative, the intrinsic second-order bias is positive, and the bias-on-bias is negative. *Lower panel:* Same as upper panel for the MAP. The second-order biases now have the same sign.

these two biases is then the second-order bias-on-bias term averaged over source ellipticity. In Figure 13 the globally-corrected estimator bias is given by the green points, and the source-averaged first-order bias by the blue dashed line, such that the green points are the difference between the blue points and the blue dashed line.

At $S/N \gtrsim 30$, both corrected and uncorrected biases are roughly within the total requirements of our fiducial survey (at least at 1σ in the bias uncertainty). Below this, neither a per-galaxy correction or a global correction succeed in reducing the bias to below this requirement, with the per-galaxy correction indistinguishable from no correction within the error bar and a global correction giving marginally lower biases. This poor performance can be understood by inspecting the upper panel of Figure 12, which tells us that higher-

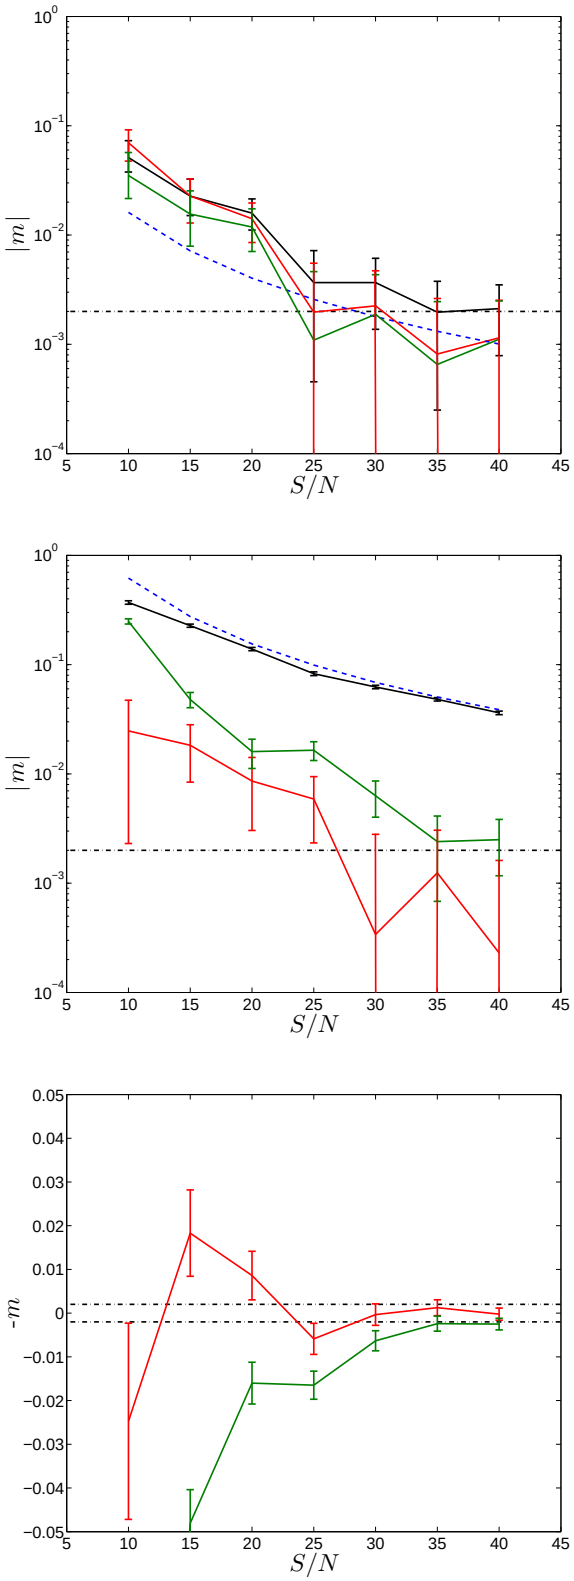

Figure 13. *Upper panel:* Uncorrected shear MLE bias from e_1 (black points), first-order bias prediction (blue dashed), first-order per-galaxy corrected residual bias (red points), globally-corrected residual bias (green points), and the total requirement for our fiducial space-based lensing survey. *Middle panel:* Same as upper panel for the MAP. *Lower panel:* Same as middle panel zooming in on low S/N and plotting only the residual biases in the bias-corrected MAP. Note the vertical axis is now on a linear scale.

order terms in the expansion of the ML equation could be important in this regime.

Considerably better performance comes however with the inclusion of a prior, as shown in the middle panel of Figure 13. The extra bias induced by the prior now means the uncorrected bias is above the requirement at all values of S/N on this plot, but now both the per-galaxy first-order correction and the global first-order correction reduce the bias by an order of magnitude to within the requirements down to $S/N \approx 30$. At higher S/N the difference between these two first-order correction schemes is given by the source-averaged bias-on-bias, which is higher order in S/N and hence small in this regime. However, below $S/N \approx 30$ this term starts to become significant (see Figure 12), and the per-galaxy correction outperforms the global correction, giving order-of-magnitude bias reduction even at $S/N = 10$. The reason for this should be clear given our previous investigation of second-order biases in the MAP - there is a cancellation between intrinsic and bias-on-bias terms, which suppresses the importance of second-order biases in the MAP. The global correction does not benefit from this cancellation, and hence suffers from large biases at low S/N. The lesson here is that higher-order terms not accounted for by a perturbative and predictive bias correction can significantly affect the performance of shear estimators at low S/N. We would expect similar conclusions to apply to the MLE/MAP as estimated from e_2.

To investigate the performance of the bias-corrected MAP on very noisy images, we repeat the image simulations with more galaxies and plot the residual biases in the first-order bias-corrected MAP in the lower panel of Figure 13. This clearly demonstrates that the per-galaxy correction brings the uncorrected bias (which is too large to be plotted on this scale) down to within the total requirement of our survey for $S/N \gtrsim 30$, and for $S/N \approx 10$ at the 1σ level defined by our error bars.

The results of this section are very encouraging, and suggest that the bias-corrected MAP not only improves biases on shape measurements over the MLE, but on shear measurements as well. Per-galaxy bias corrections seem to outperform global shear-bias corrections, due to cancellations between second-order terms that become important for $S/N \lesssim 15$. The overall importance of second-order terms is however reduced for the MAP due to the regularizing of high-$|e|$ biases, as was the case for shape measurement.

6 TWO-DIMENSIONAL SHEAR BIASES

So far our discussion has focussed on one-dimensional ellipticity and shear estimators. Whilst this permits the essential features of noise bias and bias-correction to be elucidated, it is clearly too simplistic to warrant much further study. In this section we turn our attention to the two-dimensional parameter space of ellipticity.

The multivariate generalisation of the first-order noise bias Equation (17) is derived by analogous techniques to the univariate case, the bias on parameter i being given by (Cox & Snell 1968)

$$b_i \equiv \langle \hat{\beta}_i - \beta_i \rangle = \frac{1}{2n} \bar{F}^{-1}_{ij} \bar{F}^{-1}_{km} (\bar{K}_{jkm} + 2\bar{J}_{k,jm}), \qquad (44)$$

where $\bar{F}^{-1}_{ij}$ are elements of the inverse Fisher matrix and

$$\bar{K}_{ijk} = \frac{1}{n} \sum_{m=1}^{n} \left\langle \frac{\partial^3 \log p_m}{\partial\beta_i \partial\beta_j \partial\beta_k} \right\rangle,$$
$$\bar{J}_{i,jk} = \frac{1}{n} \sum_{m=1}^{n} \left\langle \frac{\partial \log p_m}{\partial\beta_i} \frac{\partial^2 \log p_m}{\partial\beta_j \partial\beta_k} \right\rangle, \qquad (45)$$

where again we assume the expectation values are taken at the true parameter value. Note that the quantities $\bar{J}$ and $\bar{K}$ are not parameter-space tensors as they do not possess the correct transformation properties.

We will repeat our investigation of shear biases via noisy image simulations, this time finding the MLE (or MAP) for both e_1 and e_2, assuming the true (lensed) size and flux are known. To reduce the computational time, we consider galaxies having a fixed magnitude of intrinsic ellipticity $|e_s|$ of 0.3 but uniformly distributed orientation. This procedure is sometimes referred to in the literature as a 'ring test' (Nakajima & Bernstein 2007), and reduces the computational time in our case since it avoids the highly elliptical galaxies which require finer sub-pixel sampling in the convolution step. This setup also avoids the large biases in e_1 at high ellipticity. A constant shear of $(g_1, g_2) = (0.01, 0.00)$ is applied, and the shear estimate is taken to be

$$\hat{g} = \frac{\langle \hat{e} \rangle_S}{2(1 - |e_s|^2/2)}. \qquad (46)$$

The precise form of the denominator in this estimator has been chosen such that in the absence of ellipticity noise bias the above expression is an unbiased estimate of shear.

We define the multiplicative and additive shear biases in the two-dimensional case by $\langle \hat{g}_i \rangle = (\delta^K_{ij} + m_{ij}) g_j + c_i$, where δ^K_{ij} is the Kronecker delta. With the same procedure as Section 5, we find $c_1 = c_2 = m_{12} = m_{21} = 0$ by virtue of isotropy, $m_{22} = 0$ by virtue of choosing $g_2 = 0$ and

$$m_{11} = \frac{1}{2}(1 - |e_s|^2/2)^{-1} \left\langle \frac{\partial b_{e_1}(\mathbf{\Theta})}{\partial \Theta_i} \frac{\partial \Theta_i}{\partial g_1} \right\rangle_S, \qquad (47)$$

where $b_{e_1}(\mathbf{\Theta})$ now refers to the bias on e_1 as a function of the intrinsic galaxy parameters $\mathbf{\Theta}$. In Figure 14 we plot the multiplicative bias m_{11}, including the first-order per-galaxy corrected and globally corrected residual biases. The measured bias m_{21} is found to be consistent with zero. The results from this plot are encouraging - a first-order bias correction reduces the bias by an order of magnitude down to within the total requirement of our fiducial survey for $S/N \gtrsim 50$. The global and per-galaxy first-order corrections at these fairly high S/N values are indistinguishable due to the subdominance of second-order terms.

To investigate the performance of the MAP and the MLE shear biases at low S/N, we run the image simulations at $S/N = 10$ with enough galaxies to quote significant bias measurements. For the MLE at $S/N = 10$ we find an uncorrected bias of -0.026 ± 0.016, with a per-galaxy first-order correction having negligible effect and a global first-order correction reducing it to 0.010 ± 0.016. For the MAP we find that without correction, $m_{11} = -0.390 \pm 0.007$, with a per-galaxy first-order correction reducing this to -0.024 ± 0.007 and a global correction to 0.147 ± 0.007. Thus, even with a prior, a first-order correction is not sufficient to bring shear biases to within the total requirement of our fiducial survey

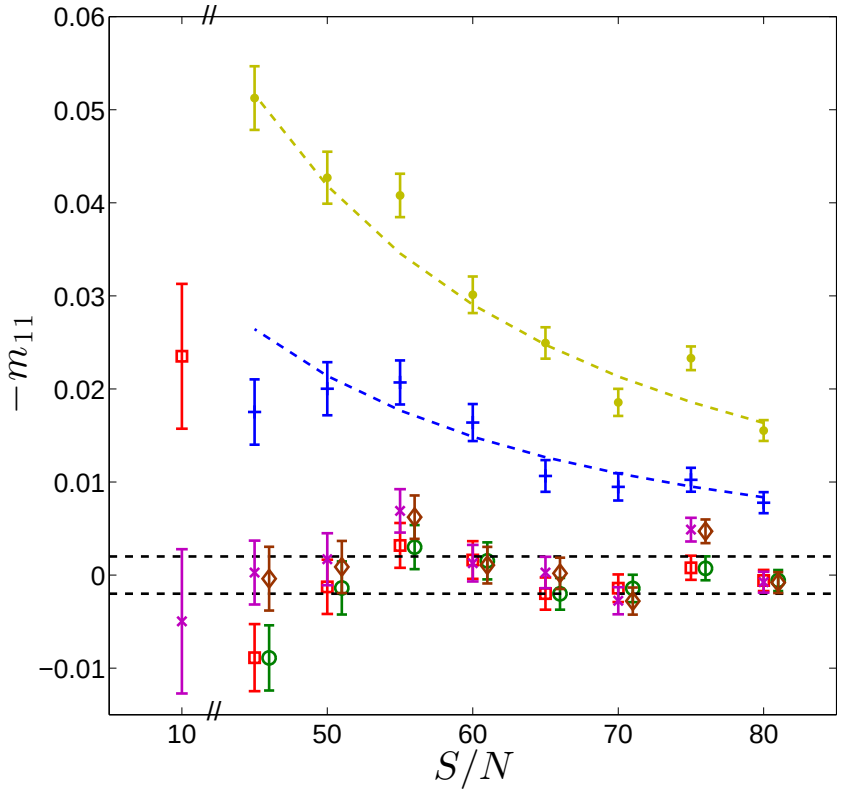

Figure 14. 2D multiplicative shear biases m_{11} for the MAP, showing uncorrected biases for a prior having $\sigma_e = 0.3$ (prior A, blue points), the first-order prediction for m_{11} with prior A (blue dashed), globally first-order corrected residual biases for prior A (green open circles), per-galaxy first-order corrected residual biases for prior A (red open squares), uncorrected biases for a prior having $\sigma_e = 0.2$ (prior B, yellow filled circles), the first-order prediction for m_{11} with prior B (yellow dashed), globally first-order corrected residual biases for prior B (brown open diamonds), and per-galaxy first-order corrected residual biases for prior B (magenta crosses). Note that residual biases are artificially offset in the horizontal direction for clarity. Note also the axis break above $S/N = 10$.

at $S/N = 10$. However, the bias is still reduced substantially by roughly an order of magnitude, with a first-order per-galaxy correction outperforming a global correction in the case of the MAP. At first glance it appears that the first-order bias-corrected MAP and the uncorrected MLE perform similarly well at $S/N = 10$, but given the large error bar on the MLE bias measurement it is difficult to draw such a conclusion reliably. Moreover, we have seen that bias correction and introduction of a prior provide significant improvement over the MLE for shape measurements at all $S/N \gtrsim 5$, and there is clearly significant improvement over uncorrected point estimates of shear at higher S/N. Indeed, it is not surprising that a first-order correction is poor at $S/N = 10$; inspection of Figure 12 shows that in the 1D case, second-order terms are comparable to first-order terms at $S/N = 10$. If this behaviour carries over to 2D, it suggests that even a second-order correction will be poor for these noisy images, as our perturbative sequence may well break down. To check this would require explicit computation of the true second-order 2D shear biases, which is beyond the scope of this work.

Finally, we investigate the sensitivity of the bias-corrected MAP to the assumed prior on intrinsic ellipticity. We repeat the ring test and bias measurements assuming the same functional form for $p(|e|)$ but setting $\sigma_e = 0.2$, which we refer to as 'prior B', with the $\sigma_e = 0.3$ distribution referred to as 'prior A'. The new value of σ_e tightens the distribution to more circular galaxies[2]. This has the effect of further down-weighting highly elliptical galaxies in the likelihood which bring with them a large noise bias to the shear estimate, at the cost of increasing the bias brought by the prior. In Figure 14 we see that the uncorrected MAP bias has increased by over a factor or 2, which can be understood by inspecting the form of the prior bias in Equation (28). For our assumed prior, the bias is roughly proportional to $1/\sigma_e^2$, which explains the increase. Despite this, our bias-correction scheme still brings the MAP bias down to levels similar to the $\sigma_e = 0.3$ case, and within the requirements of our fiducial survey. This demonstrates that the method is broadly insensitive to the prior. At $S/N = 10$ this new prior actually does significantly better, with the fiducial requirement now within the 1σ error bar of the residual bias measurement. This is clearly very encouraging, and is most likely due to the more aggressive down-weighting of high-ellipticity parts of the likelihood.

[2] Note that the true distribution is still a fixed ellipticity magntiude of 0.3 with uniform distribution in angle.

We have thus seen that the bias-corrected MAP has the capability to reduce shear biases to within the total requirements of planned surveys. At high S/N the method works very well and is broadly insensitive to the prior. At very low S/N the method appears fairly sensitive to the prior, but still has the potential to significantly reduce noise bias.

7 CONCLUSIONS AND COMPARISON TO OTHER WORK

In this work we have investigated the properties of two important point estimators for galaxy shapes - the MLE and the MAP. Our analysis has been conducted with the aid of simulated noisy images of toy galaxy models, without the effects of a PSF. Many of our results are, to an extent, predicated on these simplifying assumptions, and some have been arrived at by assuming only a one-dimensional or two-dimensional model-fitting procedure has been implemented. It is therefore legitimate to ask which of our results depend crucially on these simplifications, and which have more general applicability.

The results that should have wider generality are as follows. We have seen that noise bias arises from the non-linear functional relationship between the parameter estimate and the noise (see also Bernstein & Jarvis 2002; Hirata et al. 2004; Refregier et al. 2012; Kacprzak et al. 2012; Melchior & Viola 2012; Viola et al. 2014). Furthermore, given an analytic likelihood function and an analytic galaxy model we can predict the bias on the MLE as a perturbative series in the S/N. We have seen that these biases can become large when the galaxy size is comparable to the pixel size, as will be the case for the majority of the galaxies in future space-based weak lensing surveys. The coarseness of the pixel grid can cause a dramatic breaking of isotropy between the biases on the two ellipticity components. Using bias corrections predicted by the model itself (without fitting to external simulations), we have seen that the noise bias can be reduced by orders of magnitude down to some S/N that can only be evaluated through image simulations. The reduction of shape bias carries over to shear estimates that are composed of averaged shape estimates.

We have also seen that the large shape biases that arise at low S/N in the MLE may be mitigated by including a prior to downweight extreme best-fitting models. In this setting the MLE becomes the MAP. The extra bias incurred by including a prior may be predicted, and the MAP still possesses many of the nice asymptotic properties possessed by the MLE (Lindley 1961; Johnson 1967), although the property of invariance under general reparametrization of the model is of course lost. We have seen that the prior regularizes the distribution of best-fitting shape parameters, allowing the perturbative bias correction to be pushed to lower S/N. We found that second-order bias corrections do offer an estimate with lower bias, but the improvement is only moderate, and the additional CPU time required to compute the second-order terms precludes the significant reduction of the bias error bars that a more quantitative analysis demands.

By propagating ellipticity biases through to shear biases we have been able to directly compare our self-calibration schemes to the requirements of planned space-based weak lensing surveys. At first inspection, it seems like MAP-based shear estimates can satisfy these requirements after a bias-correction. We have seen that cancellations amongst higher-order terms resulting from our assumption of Gaussian noise can reduce these biases 'artificially'. This clearly serves as a precaution against evaluating shear estimation algorithms with the assumption of Gaussian noise - in reality the noise will not be perfectly Gaussian due to it having multiple sources. Our results should be taken as strong encouragement that these point estimators might be of considerable use in a weak lensing pipeline, and should be incorporated into algorithm testing on more realistic image simulations.

Another question that might be asked of the MLE and MAP estimates is how they compare to previous point estimators in the literature. The most relevant works to mention here are the MLE-based studies of Refregier et al. (2012) and Kacprzak et al. (2012) incorporated into the publically available code IM3SHAPE (Zuntz et al. 2013). These works cover much of the topics we have discussed here, including perturbative bias predictions, second-order bias terms, and bias-calibration. However, are approach differs to theirs in that we predict biases using the analytic perturbative biases with the MLE as a proxy for the true parameter value, rather than fitting the bias from simulations. We advocate our approach since it has much less sensitivity to the results of image simulations (which may not accurately describe the actual observed galaxy properties), which are only required to find the S/N and parameter regime where the method fails rather than being an integral part of the method per se. Similar comments could be made about other maximum likelihood methods that rely on external calibration, such as GFIT (Gentile et al. 2012).

Bernstein & Armstrong (2014) (hereafter BA14) proposed a point estimator based on a shear posterior constructed from an ensemble of galaxies subjected to a constant shear field (see Sheldon 2014; Bernstein et al. 2015, for practical implementations). The method proposed there is not directly comparable to the work here since we are concerned with estimating ellipticities and then averaging over galaxies to compute shear, rather than forming a shear posterior. Despite this difference, there are similarities to our work in that a point estimate is extracted whose bias is then computed. However, rather than studying the expansion of the likelihood around the true parameter value, BA14 expand around zero shear, assuming g is close to zero in some sense. A prior is introduced in their Equation (5) but then immediately discarded as the data are assumed to be more informative than the prior. In the limit of infinite independent data the shear posterior will become centred on the true parameter value, which need not be zero. With enough galaxies the width of the shear posterior might shrink to such an extent that the $g = 0$ point is way out in the tail of the distribution, such that both the second-order and third-order expansion of the posterior around zero shear fail to converge. Use of the MLE (or MAP) as the point estimate guarantees that the perturbative series converges in the limit of infinite independent data. It would be interesting to investigate how the MAP performs compared to the BA14 estimator. Our work also has the advantage that we do not need to assume a constant or parametrized spatially-dependent shear field as in BA14.

Unbiased estimators have also been proposed

in Madhavacheril et al. (2015) (hereafter M15). This work proposes to use the fact that certain combinations of likelihood derivatives vanish when averaged over noise realizations and evaluated at the true parameter value. By expanding these quantities around a fiducial model, one can build estimators that are unbiased at a given order in the difference between the fiducial model parameter and the truth. At leading order in this difference the 'E_1' estimator of M15 is of similar form to the leading order part of the MLE, Equation (15). The difference is that we have not had to expand around an arbitrary fiducial model. This is clearly a big advantage of the MLE - we do not need to assume anything about the true parameter value and our method does not rely on a fiducial parameter being close to the unknown true value. The leading order M15 estimate also bears some resemblance to the leading order estimator of BA14, the difference being that the BA14 is the leading-order Newton-Raphson approximation to the mode of the posterior at zero shear, whereas the M15 estimator replaces the denominator with the Fisher information and evaluates at a fiducial parameter point.

In this work we have advocated the bias-corrected MLE (or MAP) as a shape/shear estimator. The essential reason we advocate these point estimators over other frequentist methods such as re-Gaussianization (Hirata & Seljak 2003) or the KSB method (Kaiser et al. 1995) is that bias calibration can in principle be done internally rather than relying on external simulations. Furthermore, the MLE is an efficient estimator both at leading and next-to-leading order in the S/N (Rao 1961), suggesting that its mean-squared error will be better than other point estimators such as the mean of the ellipticity posterior, or a shear response weighted version of the mean as in LENSFIT (Miller et al. 2007; Kitching et al. 2008), although we have not made the requisite comparison with other methods to show this explicitly.

Finally, we note that although we have focussed on point estimators for galaxy shape and shear, in principle one could correctly propagate the probability distribution functions conditioned on the data, marginalizing over nuisance parameters to recover posteriors on the shear map, power spectra, or cosmological parameters themselves (Schneider et al. 2015; Heavens et al. 2016). This approach provides an optimal way of using all the information present in the data for inference, and is in principle preferable to using biased point estimators of galaxy shapes. However, we believe that the point estimators we have analyzed in this work are still valuable both as a shear measurement technique and as a diagnostic check on the likelihood function. Since the likelihood is a necessary ingredient of the Bayesian approach, we believe its mode could and should be used as a consistency check that the likelihood has been correctly constructed. Thus the bias-correction techniques we have advocated here fit naturally within a hierarchical Bayesian approach. In addition, the MLE and MAP are both very quick to compute, and much of the machinery required to produce shear maps from them already exists (Jarvis et al. 2015).

To conclude, we have extended the work of Refregier et al. (2012) and Kacprzak et al. (2012) to show that a simple first-order bias correction applied to the mode of the likelihood, optionally regularized by a prior, can potentially provide a point-estimate for galaxy shape capable of satisfying the stringent requirements required for percent-level constraints on dark energy. We have elucidated the origin of noise bias, and investigated the impact of the coarse grid on which galaxies will be imaged by next-generation weak lensing surveys. Our simple toy models for pixel noise and galaxy surface brightnesses have allowed us to make considerable analytic progress, allowing for some more general conclusions to be drawn regarding the performance of bias-correction on galaxy images. We have shown that ellipticity and shear biases can be reduced by orders of magnitude down to low S/N, and investigated the improvement gained by correcting for second-order biases, finding it to be small but non-negligible. We propose that these techniques be tested on the more realistic image simulations required to test shear measurement algorithms as part of current and future weak lensing experiments.

8 ACKNOWLEDGMENTS

AH is supported by a United Kingdom Space Agency *Euclid* grant, and an STFC Consolidated Grant. AT is supported by a Royal Society Wolfson Research Merit Award. We thank Giuseppe Congedo, Chris Duncan, Bryan Gillis, and Lance Miller for helpful conversations.

REFERENCES

Bartelmann M., Viola M., Melchior P., Schäfer B. M., 2012, A&A, 547, A98
Bartlett M. S., 1953, Biometrika, 40, 306
Becker M. R. et al., 2015, ArXiv e-prints 1507.05598
Bernstein G. M., Armstrong R., 2014, Mon. Not. R. Astron. Soc., 438, 1880
Bernstein G. M., Armstrong R., Krawiec C., March M. C., 2015, ArXiv e-prints 1508.05655
Bernstein G. M., Jarvis M., 2002, Astronomical Journal, 123, 583
Chang C. et al., 2013, Mon. Not. R. Astron. Soc., 428, 2695
Cox D. R., Snell E. J., 1968, Journal of the Royal Statistical Society, Series B, 30, 248
Cropper M. et al., 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8442, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 0
Gentile M., Courbin F., Meylan G., 2012, ArXiv e-prints 1211.4847
Heavens A., Alsing J., Jaffe A., Hoffmann T., Kiessling A., Wandelt B., 2016, ArXiv e-prints 1602.05345
Heymans C. et al., 2012, Mon. Not. R. Astron. Soc., 427, 146
Hirata C., Seljak U., 2003, Mon. Not. R. Astron. Soc., 343, 459
Hirata C. M. et al., 2004, Mon. Not. R. Astron. Soc., 353, 529
Jarvis M. et al., 2015, ArXiv e-prints 1507.05603
Johnson R. A., 1967, Ann. Math. Statist., 38, 1899
Kacprzak T., Zuntz J., Rowe B., Bridle S., Refregier A., Amara A., Voigt L., Hirsch M., 2012, Mon. Not. R. Astron. Soc., 427, 2711

Kaiser N., Squires G., Broadhurst T., 1995, ApJ, 449, 460
Kano Y., 1996, Journal of the Japan Statistical Society, 26, 101
Kitching T. D., Miller L., Heymans C. E., van Waerbeke L., Heavens A. F., 2008, Mon. Not. R. Astron. Soc., 390, 149
Kuijken K. et al., 2015, Mon. Not. R. Astron. Soc., 454, 3500
Laureijs R. et al., 2011, ArXiv e-prints 1110.3193
Lindley D. V., 1961, in Proc. Fourth Berkeley Symp. Math. Statist. Probab., Neyman J., ed., pp. 453–468
Madhavacheril M. S., McDonald P., Sehgal N., Slosar A., 2015, J. Cosmol. Astropart. Phys., 1, 022
Mandelbaum R. et al., 2015, Mon. Not. R. Astron. Soc., 450, 2963
Mandelbaum R. et al., 2014, ApJS, 212, 5
Mandelbaum R., Slosar A., Baldauf T., Seljak U., Hirata C. M., Nakajima R., Reyes R., Smith R. E., 2013, Mon. Not. R. Astron. Soc., 432, 1544
Massey R. et al., 2013, Mon. Not. R. Astron. Soc., 429, 661
McCullagh P., 1987, Tensor Methods in Statistics. Chapman and Hall
Melchior P., Viola M., 2012, Mon. Not. R. Astron. Soc., 424, 2757
Miller L. et al., 2013, Mon. Not. R. Astron. Soc., 429, 2858
Miller L., Kitching T. D., Heymans C., Heavens A. F., van Waerbeke L., 2007, Mon. Not. R. Astron. Soc., 382, 315
Nakajima R., Bernstein G., 2007, Ast. J., 133, 1763
Planck Collaboration et al., 2015, ArXiv e-prints 1502.01590
Rao C. R., 1961, in Berkeley Symposium on Mathematical Statistics and Probability, Vol. 1, Proc. 4th Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, pp. 531–545
Refregier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, Mon. Not. R. Astron. Soc., 425, 1951
Schneider M. D., Hogg D. W., Marshall P. J., Dawson W. A., Meyers J., Bard D. J., Lang D., 2015, ApJ, 807, 87
Schneider P., Seitz C., 1995, A&A, 294, 411
Sheldon E. S., 2014, Mon. Not. R. Astron. Soc., 444, L25
Viola M., Kitching T. D., Joachimi B., 2014, Mon. Not. R. Astron. Soc., 439, 1909
Zuntz J., Kacprzak T., Voigt L., Hirsch M., Rowe B., Bridle S., 2013, Mon. Not. R. Astron. Soc., 434, 1604